



Cisco MDS 9000 Family Configuration Guide

Cisco MDS SAN-OS Release 1.3
March, 2004

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: DOC-7816087=
Text Part Number: 78-16087-01



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0304R)

Cisco MDS 9000 Family Configuration Guide

Copyright © 2003 Cisco Systems, Inc. All rights reserved.



New and Changed Information

Table 1 summarizes the new and changed features for the *Cisco MDS 9000 Family Configuration Guide*, and tells you where they are documented. If a feature has changed in Release 1.3, a brief description of the change appears in the “Description” column, and that release is shown in the “Changed in Release” column.

Table 1 Documented Features for the *Cisco MDS 9000 Family Configuration Guide*

Feature	Description	Changed in Release	Where Documented
Running configuration information	Display Configurations based a specified feature, interface, module, or VSAN.	1.3(1)	Chapter 2, “Before You Begin”
Licensing	Access specified premium features on the switch.	1.3(1)	Chapter 3, “Obtaining and Installing Licenses”
Initial Setup Additions	Configure the full zoneset distribution and FC ID persistence features for the entire fabric during initial setup.	1.3(1)	Chapter 4, “Initial Configuration”
Automatic image synchronization	The running image is automatically synchronized in the standby supervisor module by the active supervisor module.	1.3(1)	Chapter 5, “Configuring High Availability”
Standby state	The internal standby state indicates that a switchover is possible when the redundancy state or the supervisor state display standby or HA standby.	1.3(1)	Chapter 5, “Configuring High Availability”
Terminal connection options	From the active supervisor module, you can connect to a console terminal, a Telnet terminal, or an SSH terminal.	1.3(1)	Chapter 6, “Software Images”
Standby supervisor module boot variables	The software forces the standby supervisor module to run the same version as the active supervisor module.	1.3(1)	Chapter 6, “Software Images”
Replacing modules	Ensure that the new module is running the same software version as the rest of the switch. I.	1.3(1)	Chapter 6, “Software Images”

Table 1 Documented Features for the Cisco MDS 9000 Family Configuration Guide

Feature	Description	Changed in Release	Where Documented
Transceiver and calibration information	Display real-time diagnostics information using the show interface interface-type slot/port transceiver command.	1.3(1)	Chapter 10, “Configuring Interfaces”
Buffer-to-Buffer Credit (BB_credit) display	Displays the receive and transmit BB_credit along with other pertinent interface information using the show interface bbcredit command	1.3(1)	Chapter 10, “Configuring Interfaces”
PortChannel Quiesce	Use the quiesce command on an ISL to gracefully shutdown an interface without dropping any frames.	1.3(1)	Chapter 12, “Configuring PortChannels”
Zone membership	Assign zone membership criteria is also based on the interface and domain ID, domain ID and port number, and IP address.	1.3(1)	Chapter 13, “Configuring and Managing Zones”
Inter-VSAN routing (IVR)	Access resources across VSANs without compromising other VSAN benefits.	1.3(1)	Chapter 14, “Configuring Inter-VSAN Routing”
Fabric-Device Management Interface (FDMI)	Enables management of devices using the FDMI feature.	1.3(1)	Chapter 15, “Managing FLOGI, Name Server, FDMI, and RSCN Databases”
AAA server groups	Configure remote AAA servers using server groups.	1.3(1)	Chapter 16, “Configuring Switch Security”
TACACS+ authentication	Use the Terminal Access Controller Access Control System plus (TACACS+) protocol to communicate with remote AAA servers.	1.3(1)	Chapter 16, “Configuring Switch Security”
RADIUS enhancements	Configure multiple RADIUS server groups.	1.3(1)	Chapter 16, “Configuring Switch Security”
FC-SP DHCHAP	Configure Fibre Channel Security Protocol (FC-SP) authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) provides authentication between Cisco MDS switches and other devices.	1.3(1)	Chapter 17, “Configuring Fabric Security”
FI-bre CON-nection (FICON)	Intermix FICON and Fibre Channel Protocol (FCP) traffic on the same switch without compromising scalability, availability, manageability and network security.	1.3(1)	Chapter 21, “Configuring FICON”

Table 1 Documented Features for the Cisco MDS 9000 Family Configuration Guide

Feature	Description	Changed in Release	Where Documented
Fabric Binding	Prevent unauthorized switches from joining the fabric or disrupting current fabric operations.	1.3(1)	Chapter 21, “Configuring FICON”
Registered Link Incident Report (RLIR)	Use the RLIR function to send a LIR to a registered Nx-port.	1.3(1)	Chapter 21, “Configuring FICON”
Trespass support	Use the trespass feature to enable the export of Logical Units (LUs) from the active to the passive port of a statically imported iSCSI target.	1.3(1)	Chapter 22, “Configuring IP Storage”
Internet Storage Name Service (iSNS)	Use the iSNS services to automate the discovery and management of iSCSI devices.	1.3(1)	Chapter 22, “Configuring IP Storage”
Proxy initiator	Connect all iSCSI initiators through one IPS port to make it appear as one Fibre Channel port per VSAN.	1.3(1)	Chapter 22, “Configuring IP Storage”
FCIP write accelerator	Improve application performance using the FCIP write acceleration feature.	1.3(1)	Chapter 22, “Configuring IP Storage”
FCIP compression	Allow IP packets to be compressed on the FCIP link if this feature is enabled on that link.	1.3(1)	Chapter 22, “Configuring IP Storage”
VSAN membership for iSCSI interfaces	Configure an iSCSI host to be a member of one or more VSANs.	1.3(1)	Chapter 22, “Configuring IP Storage”
Call Home enhancements	Define a Call Home destination profile, select predefined types of Call Home alerts, or filter messages based on their level of urgency.	1.3(1)	Chapter 23, “Configuring Call Home”
FC Domain ID changes	Define the default behavior to enable persistent FC IDs globally or for each VSAN.	1.3(1)	Chapter 24, “Configuring Domain Parameters”
Port rate limiting	Use the port rate limiting feature to control ingress traffic into a Fibre Channel port.	1.3(1)	Chapter 25, “Configuring Traffic Management”
Quality of Service (QoS)	Configure four priority levels for service differentiation.	1.3(1)	Chapter 25, “Configuring Traffic Management”
Auto-discovery of SCSI targets	Displays automatically discovered SCSI targets using the show scsi-target auto-poll command.	1.3(1)	Chapter 27, “Discovering SCSI Targets”
IPS SPAN source	Assign a Switched Port Analyzer (SPAN) source on the IP Storage Services (IPS) module.	1.3(1)	Chapter 28, “Monitoring Network Traffic Using SPAN”

Table 1 Documented Features for the Cisco MDS 9000 Family Configuration Guide

Feature	Description	Changed in Release	Where Documented
Per VSAN Time Out Values (TOV)	Configure different TOVs for a specified VSAN with special links like FC or IP tunnels using the ftimer command.	1.3(1)	Chapter 29, “Advanced Features and Concepts”
The zoneset import command	Imports the zoneset from the adjacent switch connected through the specified interface	1.3(2a)	Chapter 13, “Configuring and Managing Zones”
The zoneset export command	Exports the zoneset to the adjacent switch connected through the specified VSAN	1.3(2a)	
The clear system reset-reason command	Clears the reset-reason information stored in NVRAM and volatile persistent storage	1.3(2a)	Chapter 31, “Monitoring System Processes and Logs”
The update license url command	Updates an existing, expiring license file	1.3(2a)	Chapter 3, “Obtaining and Installing Licenses”
The show scsi-target pwwn command	Displays the port WWN that is assigned to each OS (Windows, AIX, Solaris, Linux, or HP-UX)	1.3(2a)	Chapter 27, “Discovering SCSI Targets”
The discover scsi-target local command	Requires the os keyword along with one of the OS options (aix , all , hpux , linux , solaris , or windows)	1.3(2a)	
Rolling upgrades	The Caching Services Module (CSM) and the IP Storage (IPS) services module use a rolling upgrade install mechanism	1.3(2a)	Chapter 6, “Software Images”
iSCSI SACK Default	The TCP SACK parameter is enabled by default for iSCSI configurations.	1.3(3)	Chapter 22, “Configuring IP Storage”
Essential Upgrade Prerequisites	Obtaining recommendations based on your current operating environment.	1.3(3)	Chapter 6, “Software Images”
iSCSI name restriction	The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.	1.3(3)	Chapter 22, “Configuring IP Storage”
Deleting directories	Deleting a specified directory deletes the entire directory and all its contents.	All	Chapter 2, “Before You Begin”

[Table 2](#) contains the history of the changes to the *Cisco MDS 9000 Family Configuration Guide*, Release 1.3. When the document is updated for the next release, these changes are incorporated into the new revision and will no longer appear in this table.

Table 2 *Documentation Changes for Cisco MDS 9000 Family Configuration Guide, Release 1.3*

Date	Description of Change	Where Changed
11/21/2003	Document created	See Table 1 .
12/19/2003	New 1.3(2a) features documented	See Table 1 .
01/13/2004	New 1.3(3) features documented	See Table 1 .



New and Changed Information iii

Preface xxxiii

Audience	xxxiii
Organization	xxxiii
Document Conventions	xxxvi
Related Documentation	xxxvi
Obtaining Documentation	xxxvii
Cisco.com	xxxvii
Documentation CD-ROM	xxxviii
Ordering Documentation	xxxviii
Documentation Feedback	xxxviii
Obtaining Technical Assistance	xxxix
Cisco.com	xxxix
Technical Support Center	xxxix
Cisco Technical Support Website	xl
Cisco Technical Support Escalation Center	xl
Obtaining Additional Publications and Information	xl

CHAPTER 1

Product Overview 1-1

Hardware Overview	1-1
Cisco MDS 9216 Fabric Switch	1-2
Cisco MDS 9500 Modular Directors	1-2
Cisco MDS 9100 Series Fixed Configuration Fabric Switches	1-3
Software Features	1-4
Licensing	1-4
High Availability	1-4
Switch Reliability	1-4
Virtual SANs	1-5
Intelligent Zoning	1-5
Inter-VSAN Routing	1-5
Trunking	1-6
PortChannels	1-6
IP Services	1-6

IP Storage	1-7
Call Home	1-7
QoS and Congestion Control	1-8
SPAN and RSPAN	1-8
Switch Management Features	1-8
Redundant Supervisor Module Management	1-9
Fabric Management	1-9
Security Management	1-9
Tools for Software Configuration	1-11
CLI	1-11
Cisco MDS 9000 Fabric Manager	1-11

CHAPTER 2

Before You Begin 2-1

About the Switch Prompt	2-2
About the CLI Command Modes	2-3
Understanding CLI Command Hierarchy	2-4
EXEC Mode Options	2-5
Configuration Mode	2-6
Configuration Mode Commands and Submodes	2-6
Navigating Through CLI Commands	2-9
Getting Help	2-9
Command Completion	2-9
Using the no and Default Forms of Commands	2-10
Entering CLI Commands	2-10
Viewing Switch Configurations	2-10
Saving a Configuration	2-13
Clearing a Configuration	2-13
Displaying Users	2-13
Sending Messages to Users	2-13
Using the ping Command	2-14
Using traceroute	2-14
Setting the Switch's Shell Timeout	2-14
Displaying VTY Sessions	2-15
Clearing VTY Sessions	2-15
Setting the Switch's Terminal Timeout	2-15
Setting the Switch's Terminal Type	2-15
Setting the Switch's Terminal Length	2-16
Setting the Switch's Terminal Width	2-16
Displaying Terminal Settings	2-16

About Flash Devices	2-17
Internal bootflash:	2-17
External CompactFlash (Slot0)	2-17
Formatting Flash Disks and File Systems	2-18
Initializing bootflash:	2-18
Formatting Slot0:	2-18
Using the File System	2-19
Setting the Current Directory	2-19
Displaying the Current Directory	2-20
Listing the Files in a Directory	2-20
Creating a New Directory	2-20
Deleting an Existing Directory	2-20
Moving Files	2-21
Copying Files	2-21
Deleting Files	2-21
Displaying File Contents	2-22
Saving Command Output to a File	2-22
Compressing and Uncompressing Files	2-22
Displaying the Last Line in a File	2-23
Executing Commands Specified in a Script	2-23
Setting the Delay Time	2-24
Role-Based CLI	2-24
Using Valid Formats and Ranges	2-25

CHAPTER 3

Obtaining and Installing Licenses	3-1
License Terminology	3-2
Licensing Model	3-4
Licensing High Availability	3-5
Options to Install a License	3-5
Obtaining a Factory-Installed License	3-6
Performing a Manual Installation	3-6
Obtaining License Key Files	3-7
Installing the License Key File	3-7
Uninstalling Licenses	3-8
Updating Licenses	3-9
License Expiry Alerts	3-10
Moving Licenses Between Switches	3-10
Displaying License Information	3-11

CHAPTER 4**Initial Configuration 4-1**

- Starting a Switch in the Cisco MDS 9000 Family 4-2
- Initial Setup Routine 4-2
 - Preparing to Configure the Switch 4-3
 - Default Login 4-3
 - Setup Options 4-4
 - Assigning Setup Information 4-5
 - Configuring Out-of-Band Management 4-5
 - In-Band Management Configuration 4-9
 - Using the setup Command 4-12
- Assigning a Switch Name 4-13
 - Assigning SNMP Switch Contact Information 4-13
- Accessing the Switch 4-14
- Where Do You Go Next? 4-14
- Verifying the Module Status 4-15
- Configuring Date and Time 4-15
 - Configuring the Time Zone 4-16
 - Setting the Daylight Saving Time Adjustment 4-16
 - NTP Configuration 4-18
 - NTP Configuration Guidelines 4-19
- Configuring the Management Port 4-20
 - Configuring Default Gateways 4-21
- Disabling a Telnet Server 4-22
- Working with Configuration Files 4-23
 - Guidelines for Creating and Using Configuration Files 4-23
 - Viewing Configuration Files 4-23
 - Downloading Configuration Files to the Switch 4-23
 - From a Remote Server 4-24
 - From an External Flash (slot0:) 4-25
 - To a Remote Server 4-25
 - To an External CompactFlash Disk 4-26
 - Saving the Configuration 4-26
- Copying Files 4-27
 - Backing up the Current Configuration 4-28
 - Rolling Back to a Previous Configuration 4-28
 - Restoring the Configured Redundancy Mode 4-29
 - Accessing Remote File Systems 4-29
 - Downgrading from a Higher Release 4-30

Deleting Files	4-30
Configuring Console Settings	4-31
Verifying the Console Configuration	4-31
Configuring COM1 Settings	4-32
Verifying the COM1 Configuration	4-32
Configuring Modem Connections	4-33
Guidelines to Configure Modems	4-33
Enabling Modem Connections	4-33
Configuring the Initialization String	4-34
Configuring the Default Initialization String	4-35
Configuring a User-Specified Initialization String	4-35
Initializing a Modem in a Powered on Switch	4-36
Configuring CDP	4-37
Clearing CDP Configurations	4-38
Displaying CDP Protocol Settings	4-39

CHAPTER 5
Configuring High Availability 5-1

About High Availability	5-2
Switchover Mechanisms	5-3
HA Switchover	5-3
Switchover Guidelines	5-3
Process Restartability	5-4
Synchronizing Supervisor Modules	5-4
Automatically Copying Images to the Standby Supervisor	5-4
Displaying HA Information	5-5

CHAPTER 6
Software Images 6-1

About Software Images	6-2
Essential Upgrade Prerequisites	6-2
Software Upgrade Mechanisms	6-4
Determining Compatibility	6-4
Performing an Automated, One-Step Upgrade	6-6
The install all command	6-6
Benefits of Using the install all Command	6-6
Recognizing Failure Cases	6-7
Using the install all Command	6-7
Sample install all Commands	6-10
Viewing the Status of an Upgrade	6-15

Performing a Manual Upgrade on a Dual Supervisor Switch	6-17
Preparing for a Manual Installation	6-17
Upgrading a Loader	6-19
Upgrading the BIOS	6-21
Upgrading Modules	6-22
Quick, One-Step Upgrade	6-23
Maintaining Supervisor Modules	6-23
Standby Supervisor's Boot Variable Version	6-23
Standby Supervisor Boot Alert	6-23
Replacing Modules	6-25
Recovering a Corrupted Bootflash	6-26
Recovery Using BIOS Setup	6-27
Recovery from the loader> Prompt	6-31
Recovery from the switch(boot)# Prompt	6-32
Recovery for Switches with Dual Supervisor Modules	6-33
Recognizing Error States	6-34
Default Factory Settings	6-35

CHAPTER 7
Managing Modules 7-1

About Modules	7-1
Supervisor Modules	7-2
Switching Modules	7-2
Verifying the Status of a Module	7-2
Viewing the State of a Module	7-3
Connecting to a Module	7-4
Reloading Modules	7-5
Reloading the Switch	7-5
Power Cycling Modules	7-5
Reloading Switching Modules	7-5
Preserving Module Configuration	7-6
Purging Module Configuration	7-7
Powering Off Switching Modules	7-7
Identifying Module LEDs	7-8
Configuring EPLDs	7-10
Displaying EPLD Versions	7-11
Default Supervisor Module Settings	7-12

CHAPTER 8**Managing System Hardware 8-1**

- Displaying Switch Hardware Inventory 8-2
- Displaying the Switch Serial Number 8-4
- Displaying Power Usage Information 8-5
- Configuring Power Supplies 8-6
 - Power Supply Guidelines 8-6
- Displaying Module Temperature 8-9
- Monitoring Fan Modules 8-10
- Monitoring Clock Modules 8-10
- Displaying Environment Information 8-11

CHAPTER 9**Configuring and Managing VSANs 9-1**

- How VSANs Work 9-2
- VSANs Versus Zones 9-4
- Default and Isolated VSANs 9-5
 - Default VSANs 9-5
 - Isolated VSANs 9-5
- VSAN Membership 9-5
- VSAN Attributes 9-6
 - Operational State of a VSAN 9-6
- Creating and Configuring VSANs 9-6
- Assigning VSAN Membership 9-7
- Deleting VSANs 9-8
- Viewing VSAN Configurations 9-9
- Default Settings 9-10

CHAPTER 10**Configuring Interfaces 10-1**

- Configuring Fibre Channel Interfaces 10-2
 - About Interface Modes 10-2
 - E Port 10-3
 - F Port 10-3
 - FL Port 10-4
 - TL Port 10-4
 - TE Port 10-4
 - SD Port 10-4
 - ST Port 10-5
 - Fx Port 10-5

B Port	10-5
Auto Mode	10-5
About Interface States	10-6
Administrative States	10-6
Operational States	10-6
Reason Codes	10-6
Configuring Fibre Channel Interfaces	10-9
Configuring a Range of Interfaces	10-9
Disabling Interfaces	10-9
Configuring Interface Modes	10-9
Configuring Administrative Speeds	10-10
Configuring Interface Descriptions	10-10
Configuring Buffer-to-Buffer Credits	10-11
Configuring Performance Buffers	10-12
Configuring Frame Encapsulation	10-12
Configuring Receive Data Field Size	10-12
Configuring the Beacon Mode	10-13
Identifying the Beacon LEDs	10-13
Configuring Switch Port Defaults	10-14
Identifying FCOT Transmitter Types	10-14
Default Settings	10-15
Configuring the Management Interface	10-16
Configuring VSAN Interfaces	10-17
Displaying Interface Information	10-17
Displaying TL Port Information	10-27
TL Port Translation Guidelines	10-29

CHAPTER 11

Configuring Trunking 11-1

About Trunking	11-1
About Trunking Protocol	11-2
Configuring Trunk Modes	11-3
Configuring Trunk-Allowed VSAN List	11-4
Trunking Configuration Guidelines	11-6
Displaying Trunking Information	11-7
Default Settings	11-8

CHAPTER 12

Configuring PortChannels 12-1

PortChannel Examples	12-2
----------------------	------

Configuring 32-port Switching Modules and Host-Optimized Ports	12-2
About PortChanneling and Trunking	12-3
About Load Balancing	12-4
Creating a PortChannel	12-5
Deleting a PortChannel	12-6
Adding Interfaces to a PortChannel	12-6
Forcing an Interface Addition	12-7
Compatibility Check	12-7
Suspended State	12-7
Deleting Interfaces from a PortChannel	12-8
Quiescing a PortChannel ISL	12-8
Considerations for PortChannel Configurations	12-9
Error Detection	12-9
Viewing PortChannel Information	12-10
Default Settings	12-12

CHAPTER 13

Configuring and Managing Zones	13-1
Zoning Features	13-2
Zoning Example	13-3
Configuring a Zone	13-4
Configuring Aliases	13-5
Zone Sets	13-6
Active and Full Zone Set Considerations	13-7
Zone Enforcement	13-9
The Default Zone	13-9
Recovering from Link Isolation	13-10
Distributing Zone Sets	13-11
Copying Zone Sets	13-11
Clearing the Zone Database	13-12
LUN Zoning	13-12
Assigning LUNs to Storage Subsystems	13-13
Read-Only Zoning	13-13
Guidelines to Configure Read-Only Zones	13-14
Configuring Read-Only Zones	13-14
Viewing Zone Information	13-15
Default Settings	13-20

Zone Implementation 13-21

CHAPTER 14

Configuring Inter-VSAN Routing 14-1

- About IVR 14-2
- IVR Features 14-2
- IVR Terminology 14-3
- IVR Guidelines 14-4
 - Domain ID Guidelines 14-4
 - Transit VSANs Guidelines 14-4
 - Border Switch Guidelines 14-5
- Configuring IVR 14-5
- Unique Domain ID Configuration Options 14-5
- Enabling IVR 14-6
- Configuring an IVR Topology 14-6
 - Creating an IVR Topology 14-6
 - Activating an IVR Topology 14-7
 - Clearing the IVR Topology 14-8
- Creating IVZs and IVZSs 14-8
 - Zones versus IVZs 14-8
 - Automatic IVZ Creation 14-9
 - Configuring and Activating IVZs and IVZSs 14-9
 - Using the force Option 14-11
- IVR Interoperability 14-12
- IVR Using LUN Zoning or Read-Only Zoning 14-12
- Clearing the IVZ Database 14-12
- Specifying IVR logging Levels 14-13
- Viewing IVR Information 14-13
- Sample Configuration 14-16

CHAPTER 15

Managing FLOGI, Name Server, FDMI, and RSCN Databases 15-1

- Displaying FLOGI Details 15-1
- Configuring the Name Server Proxy Feature 15-3
 - Registering Name Server Proxies 15-3
 - Displaying Name Server Database Entries 15-3
- Displaying FDMI 15-6
- Displaying RSCN Information 15-8
 - Sending RSCNs 15-9

Clearing RSCN Statistics 15-9

CHAPTER 16

Configuring Switch Security 16-1

Switch Management Security 16-2

CLI 16-2

SNMP Security 16-2

Switch AAA Functionalities 16-3

Authentication 16-3

Authorization 16-3

Accounting 16-4

Remote Authentication by AAA Servers 16-4

Remote Authentication Guidelines 16-4

Server Groups 16-4

AAA Service Configuration Options 16-5

Configuring RADIUS 16-6

About RADIUS 16-6

Setting the RADIUS Server Address 16-6

Setting the RADIUS Preshared Key 16-7

Setting the RADIUS Server Time-Out Interval 16-8

Setting Iterations of the RADIUS Server 16-8

Defining Vendor-Specific Attributes 16-8

VSA Format 16-8

Displaying RADIUS Server Details 16-9

Configuring TACACS+ 16-10

About TACACS+ 16-10

Advantages of TACACS+ 16-10

Enabling TACACS+ 16-11

Setting the TACACS+ Server Address 16-11

Setting the Secret Key 16-12

Setting the Timeout Value 16-12

Defining Custom Attributes for Roles 16-12

Displaying TACACS+ Server Details 16-13

Configuring Server Groups 16-14

Local AAA 16-15

No AAA Authentication 16-15

Displaying AAA Authentication 16-15

Authentication and Authorization Process 16-16

Configuring Role-Based CLI Authorization 16-18

Configuring Rules and Features for Each Role	16-18
Configuring the VSAN Policy	16-19
Displaying Role-Based CLI Information	16-20
Configuring CLI User Profiles	16-22
Creating or Updating Users	16-22
Logging out CLI Users	16-23
Displaying User Profile Information	16-23
Configuring CLI Accounting Parameters	16-24
Setting the Accounting Log Size	16-24
Displaying Accounting Configuration	16-24
Recovering Administrator Password	16-26
Configuring SSH Services	16-27
Enabling SSH Service	16-27
Generating an SSH Host Key Pair	16-27
Using the force Option	16-28
Clearing SSH Hosts	16-28
Displaying SSH Protocol Status	16-29
SNMP Security	16-30
SNMP Version 1 and Version 2c	16-30
SNMP Version 3	16-30
Restricting Switch Access	16-31
Group-Based SNMP Access	16-31
Configuring Common Roles	16-32
Creating and Modifying Users	16-32
Forcing Identical SNMP and CLI Passwords	16-33
Assigning Users to Roles	16-34
Adding or Deleting Communities	16-34
Displaying SNMP Security Information	16-35
Displaying SNMP Counter Information	16-35
Default Security Settings	16-36

CHAPTER 17

Configuring Fabric Security 17-1

About Fabric Authentication	17-2
About DHCHAP	17-3
DHCHAP Compatibility with Existing MDS Features	17-3
Configuring DHCHAP Authentication	17-3
Enabling DHCHAP	17-4
Configuring DHCHAP Authentication Modes	17-4

Configuring the DHCHAP Hash Algorithm	17-5
Configuring DHCHAP Groups	17-6
Configuring DHCHAP Passwords	17-6
Configuring Passwords for Other Devices	17-8
Configuring the DHCHAP Timeout Value	17-8
Displaying Protocol Security Information	17-9
DHCHAP AAA Authentication	17-10
Default Fabric Security Settings	17-10

CHAPTER 18

Configuring Port Security 18-1

Port Security Features	18-2
Enforcing Port Security	18-2
About Auto-Learn	18-2
Activating Port Security	18-2
Configuring Auto-Learning	18-3
Authorization Scenario	18-4
Manually Configuring Port Security	18-5
Identifying WWNs to Configure Port Security	18-5
Securing Authorized Ports	18-6
Activating the Port Security Database	18-6
Forcing Port Security Activation	18-7
Reactivating the Database	18-8
Copying the Port Security Database	18-8
Database Scenarios	18-8
Clearing the Port Security Database	18-10
Deleting the Port Security Database	18-10
Displaying Port Security Commands	18-10
Default Port Security Settings	18-12

CHAPTER 19

Configuring Fibre Channel Routing Services and Protocols 19-1

FSPF Features	19-2
FSPF Examples	19-2
Fault Tolerant Fabric	19-2
Redundant Links	19-3
Fail-over Scenarios for PortChannels and FSPF Links	19-3
Configuring FSPF Globally	19-4
Deleting the Entire FSPF Configuration	19-5

Disabling FSPF Routing Protocols	19-5
Link State Record Defaults	19-5
Configuring FSPF for a Specific Interface	19-6
Computing Route Cost	19-6
Specifying Hello Time Intervals	19-6
Specifying Dead Intervals	19-7
Disabling FSPF for Specific Interfaces	19-7
Retransmitting Intervals	19-8
Configuring Fibre Channel Routes	19-8
Clearing FSPF Counters	19-9
Broadcast Routing	19-10
In-Order Delivery	19-10
Reordering Network Frames	19-10
Reordering PortChannel Frames	19-11
Enabling In-Order Delivery	19-11
Configuring the Drop Latency Time	19-13
Displaying Latency Information	19-13
Configuring Flow Statistics	19-14
Clearing FIB Statistics	19-15
Displaying Flow Statistics	19-15
Displaying Routing and Forwarding Information	19-16
Displaying Global FSPF Information	19-18
Displaying the FSPF Database	19-18
Displaying FSPF Interfaces	19-20
Default Settings	19-20

CHAPTER 20

Configuring IP Services	20-1
Traffic Management Services	20-2
Configuring the Ethernet Management Port	20-2
Configuring the Default Gateway	20-3
Configuring the Default Network	20-4
IP Access Control Lists	20-5
IP-ACL Configuration Guidelines	20-5
Creating IP-ACLs	20-5
Adding Entries to an Existing IP-ACL	20-6
Comparing Ports	20-7
Removing Entries from an Existing IP-ACL	20-8
Applying IP-ACLs	20-8

Displaying IP-ACLs	20-10
Clearing IP-ACL Counters	20-10
Configuring IPFC	20-11
Configuring an IP Address in a VSAN	20-11
Enabling IP Routing	20-11
Configuring IP Static Routes	20-12
Viewing and Clearing ARPs	20-12
Displaying IP Interface Information	20-13
Configuring Overlay VSANs	20-14
Configuring Multiple VSANs	20-16
Configuring VRRP	20-18
VRRP Features	20-18
VRRP Functionality	20-18
Creating or Removing a Virtual Router	20-19
Enabling a Virtual Router	20-20
Adding an IP Address for a Virtual Router	20-20
Setting Priority for the Virtual Router	20-20
Setting the Time Interval for the Advertisement Packet	20-21
Preempting the Master Virtual Router	20-21
Configuring Authentication for the Virtual Router	20-22
Setting the Priority Based on Interface State	20-22
Displaying VRRP Information	20-23
Clearing VRRP Statistics	20-24
Configuring DNS Server	20-24
Displaying DNS Host Information	20-25
Default Settings	20-25

CHAPTER 21

Configuring FICON 21-1

About FICON	21-2
MDS-Specific FICON Advantages	21-2
Fabric-Optimization with VSANs	21-3
FCIP Support	21-4
PortChannel Support	21-4
VSANs for FICON and FCP Intermixing	21-4
MDS-Supported FICON Features	21-5
FICON Port Numbering	21-7
Implemented and Unimplemented Ports	21-8
Installed and Uninstalled Ports	21-8

FCIP Port Number	21-9
Port Numbering Summary	21-10
Port Addresses	21-10
FC ID Allocation	21-10
MDS FICON Prerequisites	21-10
Enabling FICON	21-11
Effects of Enabling FICON	21-11
Setting Up a Basic FICON Configuration	21-11
Manually Enabling FICON	21-14
Assigning Port Numbers to PortChannels	21-15
Assigning FCIP Interfaces	21-16
Configuring Code Page	21-16
Configuring the FC ID Last Byte	21-16
Configuring FICON Host Control	21-17
Setting the Switch State	21-17
Controlling Mainframe Access	21-17
Setting the Time Stamp	21-18
Clearing Time Stamps	21-18
Configuring FICON SNMP Control	21-18
Automatically Saving the Running Configuration	21-19
Configuring FICON Port Addresses	21-20
Blocking Ports	21-20
Prohibiting Ports	21-20
Assigning Port Address Names	21-21
FICON Configuration Files	21-22
Writing to the IPL file	21-22
Snapshot of Running Configuration	21-22
Accessing FICON Configuration Files	21-23
Applying the FICON Configuration Files	21-23
Editing FICON Configuration Files	21-23
Copying FICON Configuration Files	21-24
Port Swapping	21-25
Port Swapping Guidelines	21-25
Clearing FICON Device Allegiance	21-26
CUP Inband Management	21-26
Displaying FICON Information	21-27
Configuring Fabric Binding	21-33
Port Security versus Fabric Binding	21-33

Enforcing Fabric Binding	21-34
Enabling Fabric Binding	21-34
Configuring a List of sWWNs	21-35
Activating Fabric Binding	21-35
Forcing Fabric Binding Activation	21-36
Saving Fabric Binding Configurations	21-36
Clearing the Fabric Binding Statistics	21-37
Deleting the Fabric Binding Database	21-37
Verifying Fabric Binding Configurations	21-37
Displaying RLIR Information	21-41
Clearing RLIR Information	21-44

CHAPTER 22

Configuring IP Storage 22-1

IP Storage Services Module	22-2
Verifying the Module Status	22-3
Configuring Gigabit Ethernet Interfaces	22-4
About Gigabit Ethernet Interfaces	22-4
Basic Gigabit Ethernet Configuration	22-4
About VLANs for Gigabit Ethernet	22-5
VLAN Configuration	22-6
Interface Subnet Requirements	22-6
Managing IP Routing	22-7
Displaying the IP Route Table	22-7
Verifying Gigabit Ethernet Connectivity	22-7
Managing ARP Caches	22-8
Displaying Statistics	22-8
Displaying Gigabit Ethernet Interface Statistics	22-8
Displaying Ethernet MAC Statistics	22-9
Displaying DMA-Bridge Statistics	22-10
Displaying TCP/IP Statistics	22-10
Gigabit Ethernet High Availability	22-12
Configuring VRRP	22-12
Configuring Ethernet PortChannels	22-14
Configuring CDP	22-16
IPS Core Dumps	22-16
Configuring FCIP	22-17
About FCIP	22-17
FCIP and VE Ports	22-18
FCIP Link	22-18

FCIP Profiles	22-19
FCIP Interface	22-19
Enabling FCIP	22-20
Basic FCIP Configuration	22-20
Creating FCIP Profiles	22-20
Creating FCIP Links	22-21
Advanced FCIP Profile Configuration	22-22
Configuring TCP Listener Ports	22-22
Configuring TCP Parameters	22-23
Advanced FCIP Interface Configuration	22-28
Configuring Peers	22-28
Configuring Active Connection	22-29
Configuring the Number of TCP Connections	22-30
Enabling Time Stamps	22-30
B Port Interoperability Mode	22-32
E Port Configurations	22-34
Configuring FCIP Write Acceleration	22-35
Enabling FCIP Compression	22-36
Displaying FCIP Information	22-36
FCIP High Availability	22-39
Fibre Channel PortChannels	22-40
FSPF	22-40
VRRP	22-41
Ethernet PortChannels	22-41
Ethernet PortChannels and Fibre Channel PortChannels	22-42
Configuring iSCSI	22-43
About iSCSI	22-43
Enabling iSCSI	22-45
Routing iSCSI Requests and Responses	22-45
Presenting Fibre Channel Targets as iSCSI Targets	22-46
Dynamic Importing	22-46
Static Importing	22-47
iSCSI Virtual Target Configuration Examples	22-50
Presenting iSCSI Hosts as Virtual Fibre Channel Hosts	22-52
Dynamic Mapping	22-52
Static Mapping	22-53
Making the Dynamic Initiator WWN Mapping Static	22-55
Assigning VSAN Membership to iSCSI Hosts	22-55
Assigning VSANs to a iSCSI Interface	22-56
Configuring iSCSI Proxy Initiators	22-56

Access Control in iSCSI	22-58
Fibre Channel Zoning-Based Access Control	22-58
iSCSI-Based Access Control	22-59
Enforcing Access Control	22-60
iSCSI User Authentication	22-60
Authentication Mechanism	22-61
Advanced iSCSI Configuration	22-62
iSCSI Forwarding Mode	22-63
Displaying iSCSI Information	22-64
Displaying iSCSI Interfaces	22-64
Displaying Proxy Initiator Information	22-66
Displaying Global iSCSI Information	22-67
Displaying iSCSI Sessions	22-67
Displaying iSCSI Initiators	22-69
Displaying iSCSI Virtual Targets	22-73
Displaying IPS Statistics	22-73
Displaying iSCSI User Information	22-75
iSCSI High Availability	22-75
Multiple IPS Ports Connected to the Same IP Network	22-76
VRRP-Based High Availability	22-77
Ethernet PortChannel-Based High Availability	22-78
iSCSI Authentication Setup Guidelines	22-78
No Authentication	22-78
CHAP with Local Password Database	22-79
CHAP with External RADIUS Server	22-79
Scenario 1	22-80
Scenario 2	22-85
Configuring Storage Name Services	22-91
Creating iSNS Profiles and Tagging Profiles	22-91
Verifying iSNS Configurations	22-93
Default IP Storage Settings	22-95

CHAPTER 23

Configuring Call Home	23-1
Call Home Features	23-2
Call Home Configuration Process	23-2
Cisco AutoNotify	23-3
Configuring the Call Home Function	23-3
Assigning Contact Information	23-4
Configuring Destination Profiles	23-5

Configuring Alert Groups	23-7
Configuring Message Levels	23-8
Configuring E-Mail Options	23-9
Configuring General E-Mail Options	23-9
Configuring SMTP Server and Ports	23-9
Enabling or Disabling Call Home	23-9
Testing Call Home Communication	23-10
Displaying Call Home Information	23-10
Default Settings	23-12
Event Triggers	23-13
Call Home Message Severity Levels	23-15
Message Contents	23-16

CHAPTER 24

Configuring Domain Parameters	24-1
About fcdomain Phases	24-2
Restarting the Domain	24-3
Configuring the Domain	24-4
Setting Switch Priority	24-6
Configuring Allowed Domain ID Lists	24-6
Merging Stable Fabrics	24-7
Assigning Contiguous Domains	24-7
Disabling the fcdomain Feature	24-8
Setting the Fabric Name	24-8
Stopping Incoming RCFs	24-9
Enabling Persistent FC IDs	24-9
Configuring Persistent FC IDs Manually	24-10
Configuring Unique Area FC IDs for Some HBAs	24-11
Purging Persistent FC IDs	24-13
Displaying fcdomain Information	24-13
Default Settings	24-16

CHAPTER 25

Configuring Traffic Management	25-1
FCC	25-2
FCC Process	25-2
Enabling FCC	25-3
Assigning FCC Priority	25-3

Displaying FCC	25-3
QoS	25-4
Control Traffic	25-4
Disabling Control Traffic	25-4
Displaying Control Traffic Information	25-4
Data Traffic	25-5
Configuring Data Traffic	25-6
Enabling QoS for Data Traffic	25-6
Creating Class Maps	25-7
Defining Service Policies	25-8
Applying a Service Policy	25-8
Scheduling Traffic	25-9
Displaying Data Traffic Information	25-9
Ingress Port Rate Limiting	25-12
Default Settings	25-12

CHAPTER 26

Configuring System Message Logging	26-1
About System Message Logging	26-2
Configuring System Message Logging	26-4
Enabling Message Logging	26-4
Configuring Console Severity Level	26-5
Configuring Module Logging	26-5
Configuring Facility Severity Level	26-5
Configuring Log Files	26-6
Configuring Syslog Servers	26-6
Outgoing Syslog Server Logging Facilities	26-7
Displaying System Message Logging Information	26-8
Default Settings	26-12

CHAPTER 27

Discovering SCSI Targets	27-1
About SCSI LUN Discovery	27-1
Starting SCSI LUN Discovery	27-2
Initiating Customized Discovery	27-2
Displaying SCSI LUN Information	27-3

CHAPTER 28

Monitoring Network Traffic Using SPAN	28-1
About SPAN	28-2
SPAN Sources	28-2

IPS Source Ports	28-3
CSM Source Ports	28-3
Allowed Source Interface Types	28-3
VSAN as a SPAN Source	28-4
Guidelines to Configure VSANs as a Source	28-4
SPAN Sessions	28-5
Specifying Filters	28-5
Guidelines to Specifying Filters	28-6
SD Port Characteristics	28-6
Guidelines to Configure SPAN	28-6
Configuring SPAN	28-7
Encapsulating Frames	28-8
SPAN Conversion Behavior	28-9
Monitoring Traffic Using Fibre Channel Analyzers	28-10
Without SPAN	28-10
Using SPAN	28-11
Configuring Analyzers Using SPAN	28-11
Using a Single SD Port to Monitor Traffic	28-12
Displaying SPAN Information	28-13
Default SPAN Settings	28-14
Remote SPAN	28-15
Advantages to Using RSPAN	28-15
FC and RSPAN Tunnels	28-16
Guidelines to Configure RSPAN	28-16
ST Port Characteristics	28-17
Configuring RSPAN	28-17
Configuration in the Source Switch	28-17
Configuration in All Intermediate Switches	28-20
Configuration in the Destination Switch	28-21
Configuring An Explicit Path	28-23
Monitoring RSPAN Traffic	28-25
Sample Scenarios	28-25
Single Source with One RSPAN Tunnel	28-26
Single Source with Multiple RSPAN Tunnels	28-26
Multiple Sources with Multiple RSPAN Tunnels	28-27
Displaying RSPAN Information	28-27

CHAPTER 29

Advanced Features and Concepts 29-1

Configuring FC Timers	29-2
-----------------------	------

Configuring Timers Across All VSANs	29-2
Configuring Timers Per-VSAN	29-2
Displaying Configured FC Timer Values	29-3
Invoking the fctrace Feature	29-4
Invoking the fcping Feature	29-5
Verifying Switch Connectivity	29-6
Configuring a Fabric Analyzer	29-7
About the Cisco Fabric Analyzer	29-7
Local Text-Based Capture	29-8
Remote Capture Daemon	29-8
GUI-Based Client	29-9
Configuring the Cisco Fabric Analyzer	29-9
Capturing Frames Locally	29-9
Sending Captures to Remote IP Addresses	29-11
Clearing Configured fcanalyzer Information	29-11
Viewing Display Filters Information	29-12
Display Filters	29-12
Defining Display Filters	29-13
Displaying Filters Examples	29-13
Capture Filters	29-16
Permitted Capture Filters	29-17
Configuring World Wide Names	29-18
Configuring a Secondary MAC Address	29-18
Displaying WWN Information	29-19
Allocating Flat FC IDs	29-19
Enabling Loop Monitoring	29-20
Configuring the Switch for Interoperability	29-21
Configuring Interoperability	29-22
Verifying Interoperating Status	29-23
Using the show tech-support Command	29-27

CHAPTER 30

Configuring Fabric Configuration Servers 30-1

About FCS	30-2
Significance of FCS	30-3
Configuring FCS	30-3
Displaying FCS Information	30-4

CHAPTER 31	Monitoring System Processes and Logs	31-1
	Displaying System Processes	31-2
	Displaying System Status	31-5
	Configuring Core and Log Files	31-6
	Clearing the Core Directory	31-7
	Displaying Cores Status	31-7
	Configuring Kernel Core Dumps	31-8

INDEX



Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Organization

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	Product Overview	Presents an overview of the Cisco MDS 9000 Family of multilayer switches and directors.
Chapter 2	Before You Begin	Describes the command-line interface (CLI).
Chapter 3	Obtaining and Installing Licenses	Provides information on license types, procedure, installation, and management for the Cisco MDS SAN-OS software.
Chapter 4	Initial Configuration	Provides initial switch configuration options and switch access information.
Chapter 5	Configuring High Availability	Provides details on the high availability feature including switchover mechanisms.
Chapter 6	Software Images	Describes how to upgrade Cisco MDS 9000 Family switches, install software image files, use the Flash file system on the supervisor engine, and recover a corrupted bootflash image.
Chapter 7	Managing Modules	Explains how to display and analyze the status of each module and specifies the power on and power off process for modules.

Chapter	Title	Description
Chapter 8	Managing System Hardware	Provides details on switch hardware inventory, power usage, power supply, module temperature, fan and clock modules, and environment information.
Chapter 9	Configuring and Managing VSANs	Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs and attributes, and provides details on how to create, delete, and view VSANs.
Chapter 10	Configuring Interfaces	Explains port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces.
Chapter 11	Configuring Trunking	Explains TE ports and trunking concepts.
Chapter 12	Configuring PortChannels	Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels.
Chapter 13	Configuring and Managing Zones	Defines various zoning concepts and provides details on configuring a zone set and zone management features.
Chapter 14	Configuring Inter-VSAN Routing	Provides details on sharing resources across VSANs using the inter-VSAN Routing (IVR) feature.
Chapter 15	Managing FLOGI, Name Server, FDMI, and RSCN Databases	Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases.
Chapter 16	Configuring Switch Security	Discusses the AAA parameters, user profiles, RADIUS authentication, SSH services, and SNMP Security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options.
Chapter 17	Configuring Fabric Security	Describes the security protocols used in Cisco MDS switches to provide switch-switch and host-switch authentication for enterprise-wide fabrics.
Chapter 18	Configuring Port Security	Provides details on port security features that can prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.
Chapter 19	Configuring Fibre Channel Routing Services and Protocols	Provides details and configuration information on Fibre Channel routing services and protocols.
Chapter 20	Configuring IP Services	Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information.

Chapter	Title	Description
Chapter 21	Configuring FICON	Provides details on the Fibre Channel (FICON) interface, fabric binding, and the Registered Link Incident Report (RLIR) capabilities in Cisco MDS switches.
Chapter 22	Configuring IP Storage	Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together via IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol.
Chapter 23	Configuring Call Home	Provides details on the Call Home service and includes information on Call Home, event triggers, contact information, destination profiles, and e-mail Options.
Chapter 24	Configuring Domain Parameters	Explains the Fibre Channel domain (fcdomain) feature, which includes principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions.
Chapter 25	Configuring Traffic Management	Provides details on the quality of service (QoS) and Fibre Channel Congestion Control (FCC) features.
Chapter 26	Configuring System Message Logging	Describes how system message logging is configured and displayed.
Chapter 27	Discovering SCSI Targets	Describes how the SCSI LUN discovery feature is started and displayed.
Chapter 28	Monitoring Network Traffic Using SPAN	Describes the switched port analyzer (SPAN), identifies SPAN sources, specifies filters, explains SPAN Sessions, SD port characteristics, and configuration details.
Chapter 29	Advanced Features and Concepts	Describes the advanced configuration features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.
Chapter 30	Configuring Fabric Configuration Servers	Describes how the fabric Configuration Server (FCS) feature is configured and displayed.
Chapter 31	Monitoring System Processes and Logs	Provides information on displaying system processes and status. It also provides information on configuring core and log files, HA policy, heartbeat and watchdog checks, and upgrade resets.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
[]	Elements in square brackets are optional.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

screen font	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents:

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS SAN-OS Release 1.3(2a) and 1.3(3)*
- *Cisco MDS 9100 Series Quick Start Guide*
- *Cisco MDS 9500 Series and Cisco MDS 9216 Switch Quick Start Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9000 Family Configuration Guide*
- *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide*

- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Fabric and Device Manager User Guide*
- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family System Messages Guide*
- *Cisco MDS 9000 Family MIB Reference Guide*
- *Cisco MDS 9000 Family CIM Programming Reference Guide*

For information on VERITAS Storage Foundation™ for Networks 1.0, Cisco, refer to the following VERITAS documents available at <http://support.veritas.com/>

- *VERITAS Storage Foundation for Networks Overview*
- *VERITAS Storage Foundation for Networks Installation and Configuration Guide*
- *VERITAS Storage Foundation for Networks Obtaining and Installing Licenses*
- *VERITAS Storage Foundation for Networks GUI Administrator's Guide*
- *VERITAS Storage Foundation for Networks CLI Administrator's Guide*
- *VERITAS Storage Foundation for Networks README*

For information on IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000, refer to the following IBM documents available on the IBM TotalStorage Support web site: <http://www.ibm.com/storage/support/2062-2300/>

- Getting Started—*IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000*
- Configuration Guide—*IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000*
- Supported Hardware List—*IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000*
- Supported Software Levels—*IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000*
- Command Line Interface User's Guide—*IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000*
- Host Attachment Guide—*IBM TotalStorage SAN Volume Controller Storage Software*
- User Guide—*Subsystem Device Driver User's Guide*

Obtaining Documentation

Cisco provides several ways to obtain documentation, technical assistance, and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which may have shipped with your product. The Documentation CD-ROM is updated regularly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual or quarterly subscription.

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order monthly or quarterly subscriptions through the online Subscription Store:

<http://www.cisco.com/go/subscription>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

<http://www.cisco.com/en/US/partner/ordering/index.shtml>

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

You can e-mail your comments to mdsfeedback-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

**Note**

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Cisco provides Cisco.com, which includes the Cisco Technical Support website, as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample configurations from the Cisco Technical Support Website. Cisco.com registered users have complete access to the technical support resources on the Cisco Technical Support Website, including technical support tools and utilities.

Cisco.com

Cisco.com offers a suite of interactive, networked services that let you access Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

To obtain customized information and service, you can self-register on Cisco.com at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Technical Support Center

Cisco technical support is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available: the Cisco Technical Support Website and the Cisco Escalation Center. The type of support that you choose depends on the priority of the problem and the conditions stated in service contracts, when applicable.

We categorize Cisco technical support inquiries according to urgency:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration. There is little or no impact to your business operations.
- Priority level 3 (P3)—Operational performance of the network is impaired, but most business operations remain functional. You and Cisco are willing to commit resources during normal business hours to restore service to satisfactory levels.
- Priority level 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively impacted by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

- Priority level 1 (P1)—An existing network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Cisco Technical Support Website

The Cisco Technical Support Website provides online documents and tools to help troubleshoot and resolve technical issues with Cisco products and technologies. To access the Cisco Technical Support Website, go to this URL:

<http://www.cisco.com/techsupport>

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco Technical Support Website. Some services on the Cisco Technical Support Website require a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

<http://tools.cisco.com/RPF/register/register.do>

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco Technical Support Website, you can open a case online at this URL:

<http://www.cisco.com/techsupport/caseopen>

If you have Internet access, we recommend that you open P3 and P4 cases online so that you can fully describe the situation and attach any necessary files.

Cisco Technical Support Escalation Center

The Cisco Technical Support Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the Technical Support Escalation Center with a P1 or P2 problem, a Cisco technical support engineer automatically opens a case.

To obtain a directory of toll-free Cisco Technical Support telephone numbers for your country, go to this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Before calling, please check with your network operations center to determine the Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, please have available your service agreement number and your product serial number.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The *Cisco Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the *Cisco Product Catalog* at this URL:

http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- Cisco Press publishes a wide range of networking publications. Cisco suggests these titles for new and experienced users: *Internetworking Terms and Acronyms Dictionary*, *Internetworking Technology Handbook*, *Internetworking Troubleshooting Guide*, and the *Internetworking Design Guide*. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access *Packet* magazine at this URL:

<http://www.cisco.com/go/packet>

- iQ Magazine is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- Internet Protocol Journal is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

http://www.cisco.com/en/US/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html

- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:

http://www.cisco.com/en/US/learning/le31/learning_recommended_training_list.html



Product Overview

The Cisco MDS 9000 Family of multilayer directors and fabric switches offer intelligent fabric-switching services that realize maximum performance while ensuring high reliability levels. They combine robust and flexible hardware architecture with multiple layers of network and storage management intelligence. This powerful combination enables highly available, scalable storage networks that provide advanced security and unified management features.

The Cisco MDS 9000 Family provides intelligent networking features such as multiprotocol and multitransport integration, virtual SANs (VSANs), advanced security, sophisticated debug analysis tools, and unified SAN management.

This chapter lists the hardware features for the Cisco MDS 9000 Family and describes its software features. It includes the following sections:

- [Hardware Overview, page 1-1](#)
- [Software Features, page 1-4](#)
- [Tools for Software Configuration, page 1-11](#)

Hardware Overview

This section provides an overview of the Cisco MDS 9000 Family of multilayer directors and fabric switches.

- Cisco MDS 9216 multilayer fabric switches contain one fixed integrated supervisor module with 16 Fibre Channel ports and an expansion slot which can support up to 32 additional ports (for a total of 48 ports).
- Cisco MDS 9509 multilayer directors contain two slots for supervisor modules and 7 slots for switching or services modules providing up to 224 ports (32 ports x 7 slots).
- Cisco MDS 9506 multilayer directors contain two slots for supervisor modules and 4 slots for switching or services modules providing up to 128 ports (32 ports x 4 slots).
- Cisco MDS 9140 multilayer switches contains 40 ports (8 full rate ports, 32 host-optimized ports)
- Cisco MDS 9120 multilayer switches contains 20 ports (4 full rate ports, 16 host-optimized ports)

Cisco MDS 9216 Fabric Switch

Cisco MDS 9216 fabric switches share a consistent software architecture with the Cisco MDS 9500 Series in a semi-modular chassis. They consist of the following major hardware components:

- The chassis has two slots, one of which is reserved for the supervisor module. The supervisor module provides supervisor functions and has 16 standard, Fibre Channel ports.
- The backplane has direct plug-in connectivity to one switching module (any type).
- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to a fully loaded chassis.
- The hot-swappable fan module has four fans managing the airflow and cooling for the entire switch.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and a RS-232 (EIA/TIA-232) serial port allows switch configuration.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500m and 10km, respectively.
- The Cisco MDS 9216 supports the IP Storage Services (IPS) module. All IPS modules are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide*.

Cisco MDS 9500 Modular Directors

The Cisco MDS 9500 Series includes two multilayer, modular directors:

- The Cisco MDS 9509 Director addresses the stringent requirements of large data center storage environments and consists of the following major hardware components:
 - The chassis has nine slots, two of which are reserved for the supervisor modules.
 - Up to seven hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to seven switching modules, two integrated supervisor modules, two clock modules, and two power supplies.
 - The hot-swappable fan module has nine fans managing the airflow and cooling for the entire switch.
- The Cisco MDS 9506 Director addresses the stringent requirements of data center storage environments and consists of the following major hardware components:
 - The chassis has six slots, two of which are reserved for the supervisor modules.
 - Up to four hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to four switching modules, two integrated supervisor modules, two clock modules, and two power supplies.
 - The hot-swappable fan module has six fans managing the airflow and cooling for the entire switch.

These modular directors have the following features:

- Two redundant, hot-swappable power supplies have AC or DC connection, each of which can supply power to the entire chassis.
- Two supervisor modules ensure high availability and traffic load balancing capabilities. Each supervisor module can control the entire switch. The standby supervisor module provides redundancy in case the active supervisor module fails.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and a RS-232 serial port allows switch configuration.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500m and 10km, respectively.
- The Cisco MDS 9500 Series supports the IP Storage Services (IPS) module. All IPS modules are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide*.

Cisco MDS 9100 Series Fixed Configuration Fabric Switches

Cisco MDS 9100 Series includes two multilayer, fixed configuration (non-modular) switches:

- The Cisco MDS 9140 provides 40 ports (8 full rate ports, 32 host-optimized ports)
- The Cisco MDS 9120 is a 20 ports (4 full rate ports, 16 host-optimized ports)

These fixed configuration switches are packaged in a 1 RU enclosures and have the following features:

- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to the entire chassis.
- Two hot-swappable fan modules with two fans each manage the airflow and cooling for the entire switch.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loop (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500m and 10km, respectively.



Note

Switches in the Cisco MDS 9100 Series do not have a COM1 port (a RS-232 serial port).

Refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

Software Features

This section provides an overview of the major software features of the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Licensing

The licensing functionality is available in all switches in the Cisco MDS 9000 Family. This functionality allows you to access specified premium features on the switch after you install the appropriate license for that feature. Licenses are sold, supported, and enforced from Release 1.3(1).

See [Chapter 3, “Obtaining and Installing Licenses.”](#)

High Availability

The Cisco MDS 9500 Series of multilayer directors support application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework. The high availability (HA) software framework includes the following:

- Provides stateful redundancy for supervisor module failure by using dual supervisor modules.
- Ensures nondisruptive software upgrade capability. See [Chapter 6, “Software Images.”](#)
- Protects against link failure using the PortChannel (port aggregation) feature. See [Chapter 12, “Configuring PortChannels.”](#) This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Provides management redundancy using Virtual Router Redundancy Protocol (VRRP). See the [“Configuring VRRP” section on page 20-18](#). This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.

See [Chapter 5, “Configuring High Availability.”](#)

Switch Reliability

Switches in the Cisco MDS 9000 Family maintain internally controlled reliability services that ensure continued service with no degradation. This reliability service includes the following:

- Provides power-on self testing (POST)
- Detects errors, isolates faults, performs parity checking, and checks illegal addresses
- Enables remote diagnostics using Call Home troubleshooting features
- Displays LEDs that summarize the status of each switching module, supervisor module, power supply, and fan assembly

Virtual SANs

VSANs (virtual SANs) enable higher security and greater scalability in Fibre Channel fabrics. VSANs provide isolation among devices that are physically connected to the same fabric. VSANs allow multiple logical SANs over a common physical infrastructure. VSANs offer the following:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN thus ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical SAN. The ability to create several logical VSAN layers increases the scalability of the SAN.
- Per VSAN fabric services—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- Redundancy—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection is provided by a configured backup path between the host and the switch.
- Ease of configuration—Devices can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

See [Chapter 9, “Configuring and Managing VSANs.”](#)

Intelligent Zoning

Zoning controls access between devices in a VSAN. Zoning accomplishes the following:

- Partitions devices that use different operating systems. In a heterogeneous environment, it is often necessary to separate servers and storage devices to avoid accidental transfer of information between devices with different operating systems. Such transfers could result in corruption or deletion of data.
- Creates logical subsets of closed user groups. Closed user groups are needed to enforce security or to separate functional areas across the fabric.
- Configures groups of devices that are separate from the rest of the fabric. Based on the assigned zone membership, devices outside the zone cannot access devices internal to the zone.
- Provides temporary access between devices (zone sets). Zone restrictions can be imposed temporarily, and then restored to revert to normal operation, if desired.
- Restricts access to specific logical unit numbers (LUNs) associated with a device.
- Allows members to have only read-only access to the media within a read-only Fibre Channel zone.

See [Chapter 13, “Configuring and Managing Zones”](#) and the “VSANs Versus Zones” section on [page 9-4](#).

Inter-VSAN Routing

Using Inter-VSAN Routing (IVR), resources across VSANs can be accessed without compromising other VSAN benefits. Valuable resources like tape libraries are easily shared across VSANs without compromise. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions

See [Chapter 14, “Configuring Inter-VSAN Routing.”](#)

Trunking

Trunking is the term used to refer to an ISL link that carries one or more VSANs. Trunking ports receive and transmit Extended ISL (EISL) frames. EISL frames carry an EISL header containing VSAN information. Once EISL is enabled on an E port, that port becomes a TE port (see [Chapter 10, “Configuring Interfaces,”](#) and [Chapter 11, “Configuring Trunking”](#)). The trunking configuration is saved along with the interface information.

See the [“About PortChanneling and Trunking”](#) section on page 12-3.

PortChannels

PortChannel refers to the aggregation of multiple physical Fibre Channel ports into one logical port to provide high aggregated bandwidth, load balancing, and link redundancy. Up to 16 physical ports can be aggregated into a PortChannel. PortChannels can connect to ports across switching modules. The failure of a port in one switching module does not bring down the logical PortChannel link. Specifically, a PortChannel does the following:

- Increases the aggregate bandwidth on an ISL or EISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on a source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) that identify the flow of the frame.
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels can contain up to 16 physical links and can span multiple modules for added high availability.

See [Chapter 12, “Configuring PortChannels.”](#)

IP Services

Switches in the Cisco MDS 9000 Family support the following IP services:

- IP over Ethernet —These services are limited to management traffic.
- IP over Fibre Channel (IPFC)—IPFC (RFC 2625) specifies how IP packets are transported using encapsulation schemes. By encapsulating IP frames into Fibre Channel frames, management information is exchanged among switches without requiring a separate Ethernet connection to each switch. Each switch includes:
 - Encapsulation for IP and Address Resolution Protocol (ARP) over Fibre Channel.
 - Address resolution uses the ARP server.
- IP routing services—These services include:
 - Ethernet or TCP/IP connection.
 - Static IP routing services to enable management traffic between VSANs.
 - DNS client support.
 - The Network Time Protocol (NTP) server synchronizes the system clocks of network devices.

See [Chapter 20, “Configuring IP Services.”](#)

IP Storage

The Cisco MDS 9000 Family IP services module integrates seamlessly into the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches. Traffic can be routed between any IP storage port and any other port on a Cisco MDS 9000 Family switch. The Cisco MDS 9000 Family IP Storage Services Module supports the full range of services available on other MDS 9000 Family Switching Modules including VSANs, security, and traffic management. It uses widely known IP to cost-effectively connect to more servers and more locations over greater distances than previously possible. It delivers both Fibre Channel over IP (FCIP) and iSCSI IP storage services and is configurable on a port-by-port basis.

- FCIP highlights
 - Simplifies data protection and business continuance strategies by enabling backup, remote replication, and disaster recovery over WAN distances using open-standard FCIP tunneling.
 - Improves utilization of WAN resources for backup and replication by tunneling up to 3 virtual Inter Switch Links (ISLs) on a single Gigabit Ethernet port.
 - Reduces SAN complexity by eliminating the need to deploy and manage a separate remote connectivity platform.
 - Preserves Cisco MDS 9000 Family enhanced capabilities including VSANs, advanced traffic management, and security across remote connections.
- iSCSI highlights
 - Extends the benefits of Fibre Channel SAN-based storage to IP-enabled servers at a lower cost point than possible using Fibre Channel interconnect alone.
 - Increases storage utilization and availability through consolidation of IP and Fibre Channel block storage.
 - Transparent operation preserves the functionality of legacy storage applications such as zoning tools.
 - Extending the Benefits of Fibre Channel SANs

See [Chapter 22, “Configuring IP Storage.”](#)

Call Home

The Call Home feature detects switch failures and sends alerts along with relevant failure information. These alerts are sent through E-mail to a user-specified customer center.

See [Chapter 23, “Configuring Call Home.”](#)

QoS and Congestion Control

Switches in the Cisco MDS 9000 Family provide priority queuing and flow control services.

- The Quality of service (QoS) feature has the following advantages:
 - Provides relative bandwidth guarantee to application traffic.
 - Controls latency experienced by application traffic.
 - Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.
- Fibre Channel Congestion Control (FCC)—FCC is a flow control mechanism that alleviates congestion on Fibre Channel networks. Any switch in the network can detect congestion for an output port. The switches sample frames from the congested queue and generate messages about the congestion level upstream toward the source of the congestion. The switch closest to the source, with FCC enabled, can perform one of two actions:
 - Forwards the frames as other vendor switches do.
 - Limits the flow of frames from the port causing the congestion.

See [Chapter 25, “Configuring Traffic Management.”](#)

SPAN and RSPAN

The switched port analyzer (SPAN) feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD-port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic.

See [Chapter 28, “Monitoring Network Traffic Using SPAN”](#) and the [“Configuring a Fabric Analyzer” section on page 29-7](#)).

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch may be different from the source switch(es) provided that it is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a MDS source switch. This feature is nonintrusive and does not affect network traffic switching for any SPAN source ports.

See the [“Remote SPAN” section on page 28-15](#).

Switch Management Features

Besides the software features already listed, there are additional management features that fall into the following categories: redundant supervisor module management, fabric management, and security management

Redundant Supervisor Module Management

Series of multilayer directors support two redundant supervisor modules. They require two supervisor modules to enforce redundant supervisor module management and high availability and restartability (see [Table 1-1](#)).

Table 1-1 Supervisor Module Options in Cisco MDS 9000 Switches

Product	No. of Supervisor Modules	Slot	Features
Cisco MDS 9216	One module (includes 16 Fibre Channel ports)	Slot 1	2-slot chassis allows one optional switching module in the other slot.
Cisco MDS 9506	Two modules	Slots 5 and 6	6-slot chassis allows any switching module in the other four slots.
Cisco MDS 9509	Two modules	Slots 5 and 6	9-slot chassis allows any switching module in the other seven slots.
Cisco MDS 9140	Not applicable.		
Cisco MDS 9120	Not applicable.		

When a switch powers up and two supervisor modules are present, the module in slot 5 enters the active mode, while the second module in slot 6 enters the standby mode. All storage management functions occur on the active supervisor module. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide*.

Fabric Management

Switches in the Cisco MDS 9000 Family offer fabric management and control through the command-line interface (CLI) by using Telnet, SSH, or a serial console and through the Cisco MDS 9000 Fabric Manager tool by using the Simple Network Management Protocol (SNMP) services:

- SNMP versions 1, 2, and 3 are supported.
- Remote Monitoring (RMON) allows you to specify thresholds and monitor alarms on SNMP variables. Extended RMON alarms are available for supported Management Information Base (MIB) objects (refer to the *Cisco MDS 9000 Family MIB Reference Guide*).
- System error message logs (syslogs) are viewed through a console or Telnet session for asynchronous events such as an interface transition. Syslogs are directed to an internal log and optionally to an external server (refer to the *Cisco MDS 9000 Family System Messages Guide*).

See the “CLI” section on page 1-11, the “Cisco MDS 9000 Fabric Manager” section on page 1-11, and the “SNMP Security” section on page 16-30.

Security Management

The Cisco MDS 9000 Family of switches offer strict and secure switch management options through switch access security, port security, user authentication, and role-based access.

See [Chapter 16, “Configuring Switch Security.”](#)

Switch Access Security

Each switch can be accessed through the CLI or SNMP.

- Secure switch access—Available when you explicitly enable Secure Shell (SSH) access to the switch. SSH access provides additional controlled security by encrypting data, user IDs, and passwords. By default, Telnet access is enabled on each switch.
- SNMP access—SNMPv3 provides built-in security for secure user authentication and data encryption.
- IP Access control lists (IP-ACLs)—Provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restricts IP-related inband and out-of-band management traffic based on IP addresses (layer 3 and layer 4 information). You can use IP-ACLs to control transmissions on an interface.

Port Security

Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through syslog messages.

User Authentication

A strategy known as authentication, authorization, and accounting (AAA) is used to verify identity of, grant access, and track the actions of remote users. The Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) provide AAA solutions.

Based on the user ID and password combination provided, switches perform local authentication using a local database or remote authentication using AAA server(s). A global, preshared, secret key authenticates communication between the AAA servers. This secret key can be configured for all AAA server groups or for only a specific AAA server. This kind of authentication provides a central configuration management capability.

Role-Based Access

Role-based access control assigns roles or groups (locally through the switch or remotely using AAA servers) to users and limits access to the switch. Access is assigned based on the permission level associated with each user ID. Your administrator can provide complete access to each user or restrict access to specific read and write levels for each command.

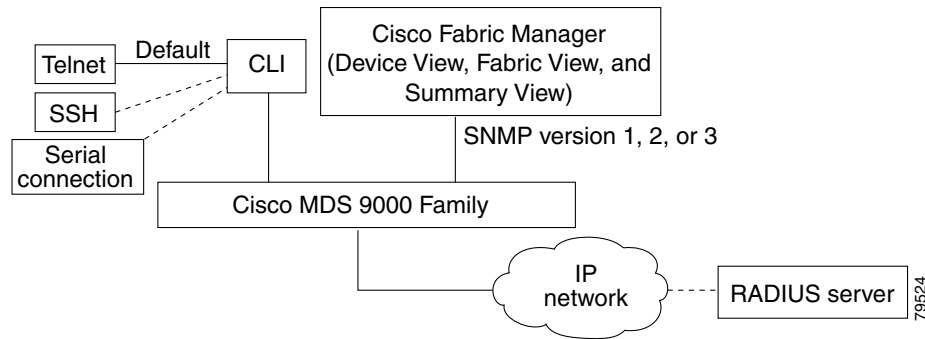
From Release 1.2(x), CLI and SNMP in all switches in the Cisco MDS 9000 Family synchronize CLI and SNMP roles. This database contains any role that is created using CLI or SNMP. You can use SNMP to modify a role that was created using CLI and vice versa. Each role in SNMP is the same as a role created or modified through the CLI.

Each role in the role database can be restricted to one or more VSAN as required.

Tools for Software Configuration

You can use one of two configuration management tools to configure your SANs: the CLI and the Cisco MDS 9000 Fabric Manager graphical user interface (see [Figure 1-1](#)).

Figure 1-1 Tools for Configuring Software



CLI

With the CLI, you can type commands at the switch prompt, and the commands are executed when you press the Enter key. The CLI parser provides command help, command completion, and keyboard sequences that allow you to access previously executed commands from the buffer history.

Cisco MDS 9000 Fabric Manager

The Cisco Fabric Manager is a Java and SNMP-based network fabric and device management tool with a GUI that displays real-time views of your network fabric and installed devices. The Cisco Fabric Manager provides three views for managing your network fabric:

- The Device View displays a continuously updated physical picture of device configuration and health conditions for a single switch.
- The Fabric View displays a view of your network fabric, including multiple switches.
- The Summary presents real-time performance statistics all active ports and channels.

The Cisco Fabric Manager provides an alternative to the CLI for most switch configuration commands. The Cisco Fabric Manager is embedded in each switch in the Cisco MDS 9000 Family.

Refer to the *Cisco MDS 9000 Fabric Manager User Guide*.



Note

Resource Manager Essentials (RME) versions 3.4 and 3.5 provide support for switches in the Cisco MDS 9000 Family. Patches are available on Cisco Connection Online (<http://www.cisco.com/>).



Before You Begin

This chapter prepares you to configure switches from the CLI. It also lists the information you need to have before you begin, and it describes the CLI command modes.

This chapter includes the following sections:

- [About the Switch Prompt, page 2-2](#)
- [About the CLI Command Modes, page 2-3](#)
- [Understanding CLI Command Hierarchy, page 2-4](#)
- [Navigating Through CLI Commands, page 2-9](#)
- [About Flash Devices, page 2-17](#)
- [Formatting Flash Disks and File Systems, page 2-18](#)
- [Using the File System, page 2-19](#)
- [Role-Based CLI, page 2-24](#)
- [Using Valid Formats and Ranges, page 2-25](#)

About the Switch Prompt

If you are connected to the console port when the switch boots up, you see the output shown in :



Note

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for installation and connection instructions.

Once the switch is powered on successfully, you see the default switch prompt (switch#). You can perform embedded CLI operations, access command history, and use command parsing functions at this prompt. The switch gathers the command string upon detecting an **Enter** (CR) and accepts commands from a terminal.

Example 2-1 Output When a Switch Boots Up

```
Auto booting bootflash:/boot-279 bootflash:/system_image;...
Booting kickstart image:bootflash:/boot-279...
.....Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
Uncompressing system image: bootflash:/system_image
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
INIT: Entering runlevel: 3

<<<<<SAN OS bootup log messages>>>>>

      ---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic configuration of
the system. Use ctrl-c to abort configuration dialog at any prompt.

Basic management setup configures only enough connectivity for
management of the system.

Would you like to enter the basic configuration dialog (yes/no): yes

<<<<<after configuration>>>>>

switch login:
```

About the CLI Command Modes

Switches in the Cisco MDS 9000 Family have two main command modes—user EXEC mode and configuration mode. The commands available to you depend on the mode you are in. To obtain a list of available commands in either mode, type a question mark (?) at the system prompt.

[Table 2-1](#) lists and describes the two commonly used modes, how to enter the modes, and the resulting system prompts. The system prompt helps you identify which mode you are in and hence, which commands are available to you.

Table 2-1 Frequently Used Switch Command Modes

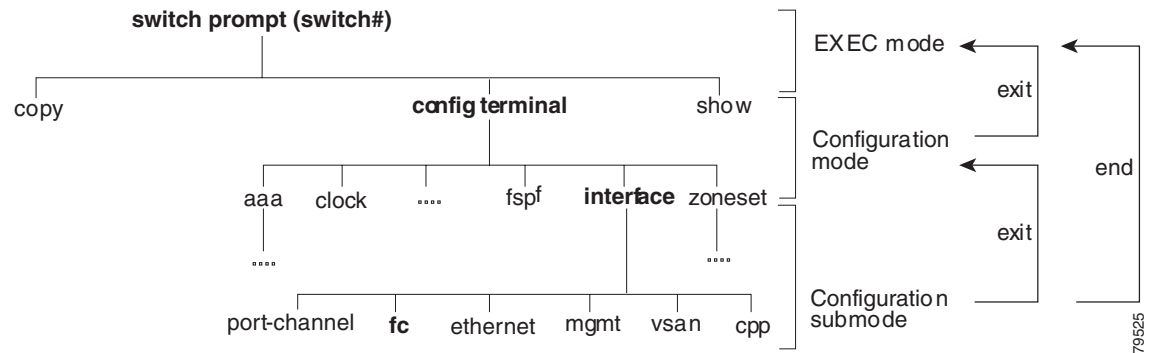
Mode	Description of Use	How to Access	Prompt
EXEC	Enables you to temporarily change terminal settings, perform basic tests, and display system information. Note Changes made in this mode are generally not saved across system resets.	At the switch prompt, enter the required EXEC mode command.	switch#
Configuration mode	Enables you to configure features that affect the system as a whole. Note Changes made in this mode are saved across system resets if you save your configuration. See the “Saving a Configuration” section on page 2-13 .	From EXEC mode, enter the config terminal command.	switch(config)#

You can abbreviate commands and keywords by entering just enough characters to make the command unique from other commands. For example, you can abbreviate the **config terminal** command to **conf t**.

Understanding CLI Command Hierarchy

The CLI commands are organized hierarchically, with commands that perform similar functions grouped under the same level. For example, all commands that display information about the system, configuration, or hardware are grouped under the **show** command, and all commands that allow you to configure the switch are grouped under the **config terminal** command. Figure 2-1 illustrates a portion of the **config terminal** command hierarchy.

Figure 2-1 CLI Command Hierarchy Example



To execute a command, you enter the command by starting at the top level of the hierarchy. For example, to configure a Fibre Channel interface, use the **config terminal** command. Once you are in configuration mode, issue the **interface** command. When you are in the interface submode, you can query the available commands there.

The following example shows how to query the available commands in the interface submode:

```

switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# interface fc1/1
switch(config-if)# ?
Interface configuration commands:
  channel-group  Add to/remove from a port-channel
  exit           Exit from this submode
  fcdomain       Enter the interface submode
  fspf           To configure FSPF related parameters
  no             Negate a command or set its defaults
  shutdown       Enable/disable an interface
  switchport     Configure switchport parameters
  
```

EXEC Mode Options

When you start a session on the switch, you begin in EXEC mode. Based on the role or group to which you belong, you have access to limited commands or to all commands (see the [“Configuring Role-Based CLI Authorization”](#) section on page 16-18). From the EXEC mode, you can enter configuration mode. Most of the EXEC commands are one-time commands, such as **show** commands, which display the current configuration status. Here is a list of EXEC mode commands:

```
switch# ?
Exec Commands:
  attach      Connect to a specific linecard
  callhome    Callhome commands
  cd          Change current directory
  clear       Reset functions
  clock       Manage the system clock
  config      Enter configuration mode
  copy        Copy from one file to another
  debug       Debugging functions
  delete      Remove files
  dir         Directory listing for files
  discover    Discover information
  exit        Exit from the EXEC
  fcping      Ping an N-Port
  fctrace     Trace the route for an N-Port.
  find        Find a file below the current directory
  format      Format disks
  install     Upgrade software
  load        Load system image
  mkdir       Create new directory
  move        Move files
  no          Disable debugging functions
  ping        Send echo messages
  purge       Deletes unused data
  pwd         View current directory
  reload      Reboot the entire box
  rmdir       Remove existing directory
  run-script  Run shell scripts
  send        Send message to all the open sessions
  setup       Run the basic SETUP command facility
  show        Show running system information
  sleep       Sleep for the specified number of seconds
  system      System management commands
  tail        Display the last part of a file
  telnet      Telnet to another system
  terminal    Set terminal line parameters
  test        Test command
  traceroute  Trace route to destination
  undebg      Disable Debugging functions (See also debug)
  write       Write current configuration
  zone        Execute Zone Server commands
```

Configuration Mode

Configuration mode allows you to make changes to the existing configuration. When you save the configuration, these commands are preserved across switch reboots. Once you are in configuration mode, you can enter interface configuration mode, zone configuration mode, and a variety of protocol-specific modes. Configuration mode is the starting point for all configuration commands. When you are in configuration mode, the switch expects configuration commands from the user.

The following example shows output from the **config terminal** command:

```
switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

Configuration Mode Commands and Submodes

The following is a list of configuration mode commands:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ?
Configure commands:
  aaa                Configure AAA
  arp                [no] remove an entry from the ARP cache
  boot               Configure boot variables
  callhome           Enter the callhome configuration mode
  clock              Configure time-of-day clock
  end                Exit from configure mode
  exit               Exit from configure mode
  fcalias             Fcalias configuration commands
  fcanalyzer          Configure cisco fabric analyzer
  fcc                Configure FC Congestion Control
  fcdomain            Enter the fcdomain configuration mode
  fcdroplatency       Configure switch or network latency
  fcflow             Configure fcflow
  fcinterop           Interop commands.
  fcns               Name server configuration
  fcroute            Configure FC routes
  fcs                Configure Fabric Config Server
  fctimer            Configure fibre channel timers
  fspf               Configure fspf
  in-order-guarantee  Set in-order delivery guarantee
  interface           Select an interface to configure
  ip                 Configure IP features
  line               Configure a terminal line
  logging             Modify message logging facilities
  no                 Negate a command or set its defaults
  ntp                NTP Configuration
  power              Configure power supply
  poweroff            Poweroff a module in the switch
  qos                Configure priority of FC control frames
  radius-server       Configure RADIUS related parameters
  role               Configure roles
  rscn                Config commands for RSCN
  snmp-server         Configure snmp server
  span               Enter SPAN configuration mode
  ssh                Configure SSH parameters
  switchname          Configure system's network name
  system              System config command
  telnet              Enable telnet
  trunk               Configure Switch wide trunk protocol
  username            Configure user information.
```


vsan	Enter the vsan configuration mode
wwn	Set secondary base MAC addr and range for additional WWNs
zone	Zone configuration commands
zoneset	Zoneset configuration commands

Configuration mode, also known as terminal configuration mode, has several submodes. Each of these submodes places you deeper in the prompt hierarchy. When you type **exit**, the switch backs out one level and returns you to the previous level. When you type **end**, the switch backs out to the user EXEC level. You can also type **Ctrl-Z** in configuration mode as an alternative to typing **end**.

**Note**

When in configuration mode, you can alternatively enter

- **Ctrl-Z** instead of the **end** command, and
- **Ctrl-G** instead of the **exit** command

You can execute an EXEC mode command from a configuration mode or submode prompt. You can issue this command from any submode within the configuration mode. When in configuration mode (or in any submode), enter the **do** command along with the required EXEC mode command. The entered command is executed at the EXEC level and the prompt resumes its current mode level.

```
switch(config)# do terminal session-timeout 0
switch(config)#
```

In this example, **terminal session-timeout** is an EXEC mode command—you are issuing an EXEC mode command using the configuration mode **do** command.

The **do** command applies to all EXEC mode commands other than the **end** and **exit** commands. You can also use the help (?) and command completion (tab) features for EXEC commands when issuing a **do** command along with the EXEC command.

Table 2-2 lists some useful command keys that can be used in both EXEC and configuration modes:

Table 2-2 Useful Command Key Description

Command	Description
Ctrl-P	Up history
Ctrl-N	Down history
Ctrl-X-H	List history
Alt-P	History search backwards Note The difference between Tab completion and Alt-P or Alt-N is that TAB completes the current word while Alt-P and Alt-N completes a previously-entered command.
Alt-N	History search forwards
Ctrl-G	Exit
Ctrl-Z	End
Ctrl-L	Clear screen

Table 2-3 displays the commonly used configuration submodes.

Table 2-3 Submodes Within the Configuration Mode

Submode Name	From Configuration Mode Enter	Submode Prompt	Configured Information
Call Home	callhome	switch(config-callhome)#	Contact, destination, and e-mail
FCS Registration	fcs register	switch(config-fcs-register)#	FCS attribute registration
	From FCS registration submode: platform name name vsan vsan-id	switch(config-fcs-register-attr-rib)#	Platform name and VSAN ID association
Fibre Channel alias	fcalias name name vsan vsan-id	switch(config-fcalias)#	Alias member
FSPF	fspf config vsan vsan-id	switch(config-(fspf-config))#	Static SPF computation, hold time, and autonomous region
Interface configuration	interface type slot/port	switch(config-if)#	Channel groups, Fibre Channel domains, FSPF parameters, switch port trunk and beacon information, and IP address
	From the VSAN or mgmt0 (management) interface configuration submode: vrrp number	switch(config-if-vrrp)#	Virtual router (see “Creating or Removing a Virtual Router” section on page 20-19)
Line console	line console	switch(config-console)#	Primary terminal console
VTY	line vty	switch(config-line)#	Virtual terminal line
Role	role name	switch(config-role)#	Rule
SPAN	span session number	switch(config-span)#	SPAN source, destination, and suspend session information
VSAN database	vsan database	switch(config-vsan-db)#	VSAN database
Zone	zone name string vsan vsan-id	switch(config-zone)#	Zone member
Zone set	zoneset name name vsan vsan-id	switch(config-zoneset)#	Zone set member

Navigating Through CLI Commands

To redisplay a command you previously entered, press the **Up Arrow** key. You can continue to press the **Up Arrow** key to see more previously issued commands. Similarly, you can press the **Down Arrow**, **Right Arrow**, **Left Arrow**, and **Delete** keys to navigate through the command history and to modify an existing command string.

Getting Help

In any command mode, you can get a list of available commands by entering a question mark (?).

```
switch# ?
```

To obtain a list of commands that begin with a particular character sequence, type in those characters followed immediately by the question mark (?). Do not include a space.

```
switch# co?
configure copy
```

To list keywords or arguments, enter a question mark in place of a keyword or argument. Include a space before the question mark. This form of help is called command syntax help, because it reminds you which keywords or arguments are applicable based on the commands, keywords, and arguments you have already entered.

```
switch# config ?
terminal Configure the system from the terminal
```



Tip

If you are having trouble entering a command, check the system prompt and enter the question mark (?) for a list of available commands. You might be in the wrong command mode or using incorrect syntax.

Command Completion

In any command mode, you can begin a particular command sequence and immediately press the **Tab** key to complete the rest of the command.

```
switch (config)# ro<Tab>
switch (config)# role <Tab>
switch (config)# role name
```

This form of help is called command completion, because it completes a word for you. If several options are available for the typed letters, all options that match those letters are presented:

```
switch(config)# fc<Tab>
fcalias          fcdomain          fcs
fcanalyzer       fcdroplacency    fcns              fctimer
fcc              fcinterop        fcroute
switch(config)# fcd<Tab>
fcdomain         fcdroplacency
switch(config)# fcd<Tab>
switch(config)# fcdomain
```

Using the no and Default Forms of Commands

You can issue the **no** form of any command to perform the following actions:

- Undo a wrongly issued command.

If you issue the **zone member** command, you can undo the results:

```
switch(config)# zone name test vsan 1
switch(config-zone)# member pwn 12:12:12:12:12:12:12:12
switch(config-zone)# no member pwn 12:12:12:12:12:12:12:12
WARNING: Zone is empty. Deleting zone test. Exit the submode.
switch(config-zone)#
```

- Delete a created facility

If you want to delete a zone that you created:

```
switch(config)# zone name test vsan 1
switch(config-zone)# exit
switch(config)# no zone name test vsan 1
switch(config)#
```

You cannot delete a zone facility called test while residing in it. You must first exit the zone submode and return to configuration mode.

Entering CLI Commands

You can configure the software in one of two ways:

- You can create the configuration for the switch interactively by issuing commands at the CLI prompt.
- You can create an ASCII file containing a switch configuration and then load this file on the required system. You can then use the CLI to edit and activate the file (see the [“Working with Configuration Files”](#) section on page 4-23).

Viewing Switch Configurations

You can view the ASCII form of the configuration file when required. To view the current configuration tree from the EXEC prompt, issue the **show running-config** command. If the running configuration is different from the startup configuration, issue the **show startup-config** command to view the ASCII version of the current startup configuration that was used to boot the switch.

You can also gather specific information on the entire switch configuration by issuing the relevant **show** commands. Configurations are displayed based a specified feature, interface, module, or VSAN. Available **show** commands for each feature are briefly described in this section and listed at the end of each chapter.

Examples 2-2 to 2-8 display a few **show** command examples.

Example 2-2 Displays Details on the Specified Interface

```
switch# show interface fc1/1
fc1/1 is up
  Hardware is Fibre Channel, 20:01:ac:16:5e:4a:00:00
  vsan is 1
  Port mode is E
```

```

Speed is 1 Gbps
Beacon is turned off
FCID is 0x0b0100
  0 frames input, 0 bytes, 0 discards
  0 runts, 0 jabber, 0 too long, 0 too short
  0 input errors, 0 CRC, 0 invalid transmission words
  0 address id, 0 delimiter
  0 EOF abort, 0 fragmented, 0 unknown class
  0 frames output, 0 bytes, 0 discards
Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits

```

Example 2-3 Displays the Software and Hardware Version

```

switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:      version 1.0.8
  loader:    version 1.1(2)
  kickstart: version 2.0(1) [build 2.0(0.6)] [gdb]
  system:    version 2.0(1) [build 2.0(0.6)] [gdb]

  BIOS compile time:      08/07/03
  kickstart image file is: bootflash:///m9500-sflek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time: 10/25/2010 12:00:00
  system image file is:   bootflash:///m9500-sflek9-mzg.2.0.0.6.bin
  system compile time:    10/25/2020 12:00:00

Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:      0 blocks (block size 512b)

  172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

  Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
  Reason: Reset Requested by CLI command reload
  System version: 2.0(0.6)
  Service:

```

Example 2-4 Displays the Running Configuration

```

switch# show running
Building Configuration ...
  interface fc1/1
  interface fc1/2
  interface fc1/3
  interface fc1/4
  interface mgmt0
ip address 172.22.95.112 255.255.255.0
no shutdown
vsan database
boot system bootflash:system-237; sup-1

```

```
boot kickstart bootflash:boot-237 sup-1
callhome
ip default-gateway 172.22.95.1
switchname switch
trunk protocol enable
username admin password 5 /AFDAMD4B2xK2 role network-admin
```

Example 2-5 *Displays the Difference between the Running and Startup Configuration*

```
switch# show running diff
Building Configuration ...
*** Startup-config
--- Running-config
***** 1,16 ****
    fcip enable
    ip default-gateway 172.22.91.1
    iscsi authentication none
    iscsi enable
! iscsi import target fc
    iscsi virtual-target name vt
        pWWN 21:00:00:04:cf:4c:52:c1
    all-initiator-permit
--- 1,20 ----
    fcip enable
+ aaa accounting logsize 500
+
+
+
    ip default-gateway 172.22.91.1
    iscsi authentication none
    iscsi enable
! iscsi initiator name junk
    iscsi virtual-target name vt
        pWWN 21:00:00:04:cf:4c:52:c1
    all-initiator-permit
```

Example 2-6 *Displays the Configuration for a Specified Interface*

```
switch# show running interface fc2/9
interface fc2/9
switchport mode E
no shutdown
```



Note

The **show running interface** command is different from the **show interface** command.

Example 2-7 *Displays the Configuration for all Interfaces in a 16-Port Module*

```
switch# show running interface fc2/10 - 12
interface fc2/10
switchport mode E
no shutdown

interface fc2/11
switchport mode E
no shutdown

interface fc2/12
switchport mode FL
no shutdown
```

Example 2-8 Displays the Configuration Per VSAN

```
switch# show running vsan 1
Building Configuration ...
zone name m vsan 1
  member pwn 21:00:00:20:37:60:42:5c
  member pwn 21:00:00:20:37:4b:00:a2
zoneset name m vsan 1
  member m
zoneset activate name m vsan 1
```

Saving a Configuration

To save the configuration, enter the **copy running-config startup-config** command from the EXEC mode prompt to save the new configuration into nonvolatile storage. Once this command is issued, the running and the startup copies of the configuration are identical.

See the “Copying Files” section on page 4-27.

Clearing a Configuration

To clear a startup configuration, enter the **write erase** command from the EXEC mode prompt. Once this command is issued, the switch’s startup configuration reverts to factory defaults. The running configuration is not affected. The **write erase** command erases the entire startup configuration with the exception of any configuration that affects the loader functionality.

The **write erase boot** command only erases the configuration that affects the loader functionality. The loader functionality configuration includes the boot variables and the mgmt0 IP configuration information (IP address, netmask and default gateway).

```
switch# write erase boot
This command will erase the boot variables and the ip configuration of interface mgmt 0
```

Displaying Users

The **show users** command displays all users currently accessing the switch.

```
switch# show users
admin pts/7 Jan 12 20:56 (10.77.202.149)
admin pts/9 Jan 12 23:29 (modena.cisco.com)
admin pts/11 Jan 13 01:53 (dhcp-171-71-49-49.cisco.com)
```

Sending Messages to Users

The **send** command sends a message to all active CLI users currently using the switch. This message is restricted to 80 alphanumeric characters with spaces.

This example sends a warning message to all active users about the switch being shut down.

```
switch# send Shutting down the system in 2 minutes. Please log off.

Broadcast Message from admin@excal-112
(/dev/pts/3) at 16:50 ...
Shutting down the system in 2 minutes. Please log off.
```

Using the ping Command

The **ping** command verifies the connectivity of a remote host or server by sending echo messages.

The syntax for this command is **ping** *<host or ip address>*

```
switch# ping 171.71.181.19
PING 171.71.181.19 (171.71.181.19): 56 data bytes
64 bytes from 171.71.181.19: icmp_seq=0 ttl=121 time=0.8 ms
64 bytes from 171.71.181.19: icmp_seq=1 ttl=121 time=0.8 ms

--- 171.71.181.19 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0.8/0.8/0.8 ms
```

To abnormally terminate a ping session, type the **Ctrl-C** escape sequence

Using traceroute

The **traceroute** command prints the routes taken by a specified host or IP address.

The syntax for this command is **traceroute** *<host or ip address>*

```
switch# traceroute www.cisco.com
traceroute to www.cisco.com (171.71.181.19), 30 hops max, 38 byte packets
 1 kingfisher1-92.cisco.com (172.22.92.2) 0.598 ms 0.470 ms 0.484 ms
 2 nbulab-gw1-bldg6.cisco.com (171.71.20.130) 0.698 ms 0.452 ms 0.481 ms
 3 172.24.109.185 (172.24.109.185) 0.478 ms 0.459 ms 0.484 ms
 4 sjc12-lab4-gw2.cisco.com (172.24.111.213) 0.529 ms 0.577 ms 0.480 ms
 5 sjc5-sbb4-gw1.cisco.com (171.71.241.174) 0.521 ms 0.495 ms 0.604 ms
 6 sjc12-dc2-gw2.cisco.com (171.71.241.230) 0.521 ms 0.614 ms 0.479 ms
 7 sjc12-dc2-cec-css1.cisco.com (171.71.181.5) 2.612 ms 2.093 ms 2.118 ms
 8 www.cisco.com (171.71.181.19) 2.496 ms * 2.135 ms
```

To abnormally terminate a traceroute session, enter **Ctrl-C**.

Setting the Switch's Shell Timeout

Use the **exec-timeout** command in configuration mode to configure the lifetime of all terminal sessions on that switch. When the time limit configured by this command is exceeded, the shell exits and closes that session. The syntax for this command from is **exec-timeout** *minutes*

The default is 30 minutes. You can configure different timeout values for a console or a virtual terminal line (VTY) session. You can set the **exec-timeout** value to 0 to disable this feature so the session remains active until you exit the switch. This change is saved in the configuration file.

- From the console:

```
switch(config)# line console
switch(config-console)# exec-timeout 60
```

Specifies the current console shell timeout to be 60 minutes.

- From a VTY session (Telnet or SSH):

```
switch(config)# line vty
switch(config-line)# exec-timeout 60
```

Specifies the current console shell timeout to be 60 minutes.

Displaying VTY Sessions

Use the **show line** command to display all configured VTY sessions:

```
switch# show line
line Console:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:       1 bit(s)
  Parity:         none
  Modem In: Disable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Statistics: tx:5558511 rx:5033958 Register Bits:RTS|CTS|DTR|DSR|CD|RI
line Aux:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:       1 bit(s)
  Parity:         none
  Modem In: Disable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Hardware Flowcontrol: ON
  Statistics: tx:35 rx:0 Register Bits:RTS|DTR
```

Clearing VTY Sessions

Use the **clear line** command to close a specified VTY session:

```
switch# clear line Aux
```

Setting the Switch's Terminal Timeout

Use the **terminal session-timeout** command in EXEC mode to configure the automatic logout time for the current terminal session on that switch. When the time limit configured by this command is exceeded, the switch closes that session and exits.

The syntax for this command from is **terminal session-timeout** *minutes*

The default is 30 minutes. You can set the **terminal session-timeout** value to 0 to disable this feature so the terminal remains active until you choose to exit the switch. This change is not saved in the configuration file.

```
switch# terminal session-timeout 600
```

Specifies the terminal timeout to be 600 minutes for the current session.

Setting the Switch's Terminal Type

Use the **terminal terminal-type** command in EXEC mode to specify the terminal type for a switch:

The syntax for this command is **terminal terminal-type** *terminal-type*

```
switch# terminal terminal-type vt100
```

Specifies the terminal type. The *terminal-type* string is restricted to 80 characters and must be a valid type (for example vt100 or xterm). If a Telnet or SSH session specifies an unknown terminal type, the switch uses the vt100 terminal by default.

Setting the Switch's Terminal Length

To set the terminal screen length for the current session, use the **terminal length** command in EXEC mode. This command is specific to only the console port. Telnet and SSH sessions set the length automatically.

The syntax for this command is **terminal length** *lines*

```
switch# terminal length 20
```

Sets the screen length for the current session to 20 lines for the current terminal session. The default is 24 lines.

Setting the Switch's Terminal Width

To set the terminal screen width for the current session, use the **terminal width** command in EXEC mode. This command is specific to only the console port. Telnet and SSH sessions set the width automatically.

The syntax for this command is **terminal width** *columns*

```
switch# terminal width 86
```

Sets the screen length for the current session to 86 columns for the current terminal session. The default is 80 columns.

Displaying Terminal Settings

The show terminal command displays the terminal settings for the current session:

```
switch# show terminal
TTY: Type: "vt100"
Length: 24 lines, Width: 80 columns
Session Timeout: 525600 minutes
```

About Flash Devices

Every switch in the Cisco MDS 9000 Family contains one internal bootflash (see [Figure 2-2](#)). The Cisco MDS 9500 Series additionally contains one external CompactFlash called slot0 (see [Figure 2-2](#) and [Figure 2-3](#)).

Figure 2-2 Flash Devices in the Cisco MDS 9000 Supervisor Module

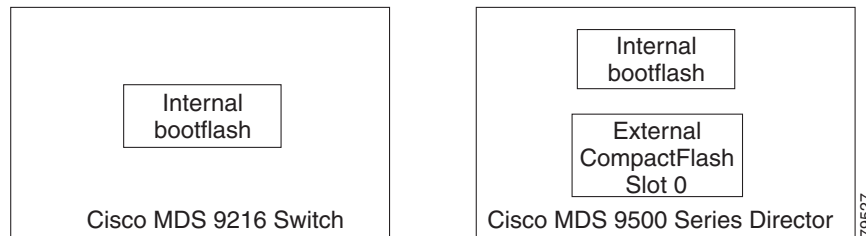
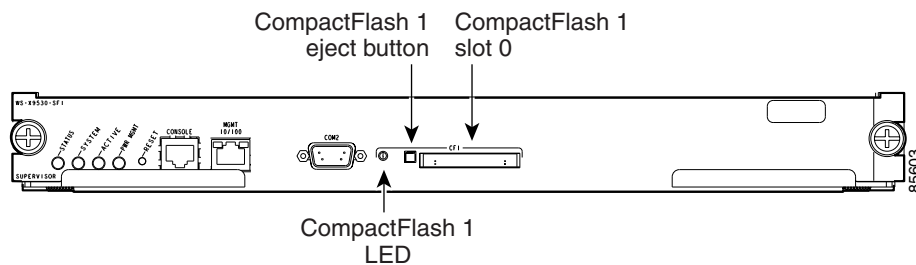


Figure 2-3 External CompactFlash in the Cisco MDS 9000 Supervisor Module



Internal bootflash:

All switches in the Cisco MDS 9000 Family have one internal bootflash: that resides in the supervisor or switching module. You have access to two directories within the internal bootflash: file system.

- The volatile: directory which provides temporary storage, and is also the default. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash (nonvolatile storage): directory which provides permanent storage. The files in bootflash are preserved through reboots and power outages.

External CompactFlash (Slot0)

Cisco MDS 9500 Series directors contain an additional external CompactFlash called slot0:

The external CompactFlash, an optional device for MDS 9500 Series directors, can be used for storing software images, logs, and core dumps.

Formatting Flash Disks and File Systems

By formatting a flash disk or a file system, you are essentially clearing out the contents of the disk or the file system and restoring it to its factory-shipped state (see the [“About Flash Devices”](#) section on page 2-17 and [“Using the File System”](#) section on page 2-19 for additional information).

Initializing bootflash:

When a switch is shipped, the **init system** command is already performed and you do not need to issue it again. Initializing the switch resets the entire internal disk and erases all data in the bootflash: partition. The internal disk is composed of several file systems with bootflash: being one of them. All files in bootflash: are erased and you must download the system and kickstart images again. After issuing an **init system** command, you don't have to format the bootflash: again since bootflash: is automatically formatted.

**Note**

The **init system** command also installs a new loader from the existing (running) kickstart image. You can access this command from the `switch(boot) #` prompt (see [Chapter 6, “Software Images”](#)).

If bootflash: is found corrupted during a boot sequence, you will see the following message:

```
ERROR:bootflash: has unrecoverable error; please do "format bootflash:"
```

Use the **format bootflash:** command to only format the bootflash: file system. You can issue the **format bootflash:** command from either the `switch#` or the `switch(boot) #` prompts.

If you issue the **format bootflash:** command, you must download the kickstart and system images again.

Formatting Slot0:

Be sure to format an external CompactFlash device before using it to save files or images.

You can verify if the external CompactFlash device is formatted by inserting it into slot0: and issuing the **dir slot0:** command.

- If the external CompactFlash device is already formatted, you can see file system usage information (along with any existing files).
- If the external CompactFlash device is unformatted (corrupted), you will see the following message:

```
Device unavailable
```

In this case, you need to format the CompactFlash device using the **format slot0:** command.

**Note**

The slot0: file system cannot be accessed from the standby the `loader>` prompt or the `switch(boot) #` prompt, if the disk is inserted after booting the switch.

**Caution**

The SAN-OS software only supports Cisco-certified CompactFlash devices that are formatted using Cisco MDS switches. Using uncertified CompactFlash devices may result in unpredictable consequences; formatting CompactFlash devices using other platforms may result in errors.

Using the File System

The switch provides the following useful functions to help you manage software image files and configuration files:

- [Setting the Current Directory, page 2-19](#)
- [Displaying the Current Directory, page 2-20](#)
- [Listing the Files in a Directory, page 2-20](#)
- [Creating a New Directory, page 2-20](#)
- [Deleting an Existing Directory, page 2-20](#)
- [Moving Files, page 2-21](#)
- [Copying Files, page 2-21](#)
- [Deleting Files, page 2-21](#)
- [Displaying File Contents, page 2-22](#)
- [Saving Command Output to a File, page 2-22](#)
- [Compressing and Uncompressing Files, page 2-22](#)
- [Displaying the Last Line in a File, page 2-23](#)
- [Executing Commands Specified in a Script, page 2-23](#)
- [Setting the Delay Time, page 2-24](#)

Setting the Current Directory

The **cd** command changes the current directory level to a specified directory level. CLI defaults to the volatile: files system. This command expects a directory name input.

**Tip**

Any file saved in the volatile: file system will be erased when the switch reboots.

The syntax for this command is **cd** *directory name*

This example changes the current directory to the mystorage directory that resides in the slot0 directory:

```
switch# cd slot0:mystorage
```

This example changes the current directory to the mystorage directory that resides in the current directory.

```
switch# cd mystorage
```

If the current directory is slot0:mydir, this command changes the current directory to slot0:mydir/mystorage.

Displaying the Current Directory

The **pwd** command displays the current directory location. This example changes the directory and displays the current directory.

```
switch# cd bootflash:
switch# pwd
bootflash:
```

Listing the Files in a Directory

The **dir** command displays the contents of the current directory or the specified directory. The syntax for this command is **dir** *directory or file name*

This example shows how to list the files on the default volatile: file system:

```
switch# dir
      Usage for volatile: filesystem
                0 bytes total used
      20971520 bytes free
      20971520 bytes available
```

Creating a New Directory

The **mkdir** command creates a directory at the current directory level or at a specified directory level.

The syntax for this command is **mkdir** *directory name*

This example creates a directory called test in the slot0 directory.

```
switch# mkdir slot0:test
```

This example creates a directory called test at the current directory level.

```
switch# mkdir test
```

If the current directory is slot0:mydir, this command creates a directory called slot0:mydir/test.

Deleting an Existing Directory

The **rmdir** command deletes an existing directory at the current directory level or at a specified directory level. The directory must be empty to be deleted.

The syntax for this command is **rmdir** *directory name*

This example deletes the directory called test in the slot0 directory.

```
switch# rmdir slot0:test
```

This example deletes the directory called test at the current directory level.

```
switch# rmdir test
```

If the current directory is slot0:mydir, this command deletes the slot0:mydir/test directory.

Moving Files

The **move** command removes a file from the source directory and places it in the destination directory. If a file with the same name already exists in the destination directory, that file is overwritten by the moved file.

This example moves the file called samplefile from the slot0 directory to the mystorage directory.

```
switch# move slot0:samplefile slot0:mystorage/samplefile
```

This example moves a file from the current directory level.

```
switch# move samplefile mystorage/samplefile
```

If the current directory is slot0:mydir, this command moves slot0:mydir/samplefile to slot0:mydir/mystorage/samplefile.

Copying Files

The **copy** command copies a file.

This example copies the file called samplefile from the external CompactFlash (slot0) directory to the mystorage directory.

```
switch# copy slot0:samplefile slot0:mystorage/samplefile
```

This example copies a file from the current directory level.

```
switch# copy samplefile mystorage/samplefile
```

If the current directory is slot0:mydir, this command copies slot0:mydir/samplefile to slot0:mydir/mystorage/samplefile.

You can also use the **copy** command to upload and download files from the slot0: or bootflash: file system to or from a FTP, TFTP, SFTP, or SCP server (see the [“Copying Files” section on page 4-27](#)).

Deleting Files

The **delete** command deletes a specified file or the specified directory and all its contents (see the [“Deleting Files” section on page 4-30](#)).

This example shows how to delete a file from the bootflash: directory (assuming you are already in the bootflash: directory):

```
switch# delete dns_config.cfg
```

This example shows how to delete a file from an external CompactFlash (slot0):

```
switch# delete slot0:dns_config.cfg
```

This example deletes the entire my-dir directory and all its contents:

```
switch# delete bootflash:my-dir
```



Caution

If you specify a directory, the **delete** command deletes the entire directory and all its contents.

Displaying File Contents

The **show file** command displays the contents of a specified file in the file system.

The syntax for this command is **show file *file_name***

This example displays the contents of the test file that resides in the slot0 directory.

```
switch# show file slot0:test
config t
Int fc1/1
no shut
end
show int
```

This example displays the contents of a file residing in the current directory.

```
switch# show file myfile
```

Saving Command Output to a File

You can force all screen output to go to a file by appending **> filename** to any command. For example, enter **show interface > samplefile** at the EXEC mode switch prompt to save the interface configuration to *samplefile*—a file created at the same directory level. At the EXEC mode switch prompt, issue a **dir** command to view all files in this directory, including the recently saved *samplefile*. See [Chapter 4, “Initial Configuration,”](#) for information on saving and copying configuration files, and [Chapter 6, “Software Images,”](#) for information on saving and copying software images.



Note

Redirection is allowed only if the current directory is on the `volatile:` (default) or `slot0:` file systems. Redirection is not allowed if the current directory is on the `bootflash:` file system. The current directory can be viewed using the **pwd** command and changed using the **cd** command.

Compressing and Uncompressing Files

The **gzip** command compresses (zips) the specified file using LZ77 coding.

This example directs the output of the **show tech-support** command to a file (Samplefile) and then zips the file and displays the difference in the space used up in the `volatile:` directory:

```
switch# show tech-support > Samplefile
Building Configuration ...
switch# dir
    1525859      Jul 04 00:51:03 2003 Samplefile
Usage for volatile://
    1527808 bytes used
    19443712 bytes free
    20971520 bytes total
switch# gzip volatile:Samplefile
switch# dir
    266069      Jul 04 00:51:03 2003 Samplefile.gz
Usage for volatile://
    266240 bytes used
    20705280 bytes free
    20971520 bytes total
```


The **gunzip** command uncompresses (unzips) LZ77 coded files.

This example unzips the file that was compressed in the previous example:

```
switch# gunzip samplefile
/volatile/samplefile.gz: No such file or directory
switch# gunzip Samplefile
switch# dir
    1525859      Jul 04 00:51:03 2003 Samplefile
Usage for volatile://
    1527808 bytes used
    19443712 bytes free
    20971520 bytes total
```

Displaying the Last Line in a File

The **tail** command displays the last lines (tail end) of a specified file.

The syntax for this command is **tail** *<file name>* [*<number of lines>*]

```
switch# tail mylog 10
```

You see the last 10 lines of the mylog file.

Executing Commands Specified in a Script

The **run-script** command executes the commands specified in a file. To use this command, be sure to create the file and specify commands in the required order.



Note

You cannot create the script files at the switch prompt. You can create the script file on an external machine and copy it to the bootflash: directory. This section assumes that the script file resides in the bootflash: directory.

The syntax for this command is **run-script** *file_name*

This example displays the CLI commands specified in the testfile that resides in the slot0 directory.

```
switch# show file slot0:testfile
conf t
interface fc 1/1
no shutdown
end
sh interface fc1/1
```

This file output is in response to the **run-script** command executing the contents in the testfile file:

```
switch# run-script slot0:testfile
'conf t'
Enter configuration commands, one per line. End with CNTL/Z.

'interface fc1/1'

'no shutdown'

'end'

'sh interface fc1/1'
fc1/1 is down (Fcot not present)
```

```

Hardware is Fibre Channel
Port WWN is 20:01:00:05:30:00:48:9e
Admin port mode is auto, trunk mode is on
vsan is 1
Beacon is turned off
Counter Values (current):
  0 frames input, 0 bytes, 0 discards
  0 runts, 0 jabber, 0 too long, 0 too short
  0 input errors, 0 CRC, 0 invalid transmission words
  0 address id, 0 delimiter
  0 EOF abort, 0 fragmented, 0 unknown class
  0 frames output, 0 bytes, 0 discards
  Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
  Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
Counter Values (5 minute averages):
  0 frames input, 0 bytes, 0 discards
  0 runts, 0 jabber, 0 too long, 0 too short
  0 input errors, 0 CRC, 0 invalid transmission words
  0 address id, 0 delimiter
  0 EOF abort, 0 fragmented, 0 unknown class
  0 frames output, 0 bytes, 0 discards
  Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
  Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits

```

Setting the Delay Time

The **sleep** command delays an action by a specified number of seconds.

The syntax for this command is **sleep <seconds>**

```
switch# sleep 30
```

You will see the switch prompt return after 30 seconds.

This command is useful within scripts. For example, if you create a script called test-script:

```

switch# show file slot0:test-script
discover scsi-target remote
sleep 10
show scsi-target disk

switch# run-script slot0:test-script

```

When you execute the slot0:test-script, the switch software executes the **discover scsi-target remote** command, and then waits for 10 seconds before executing the **show scsi-target disk** command.

Role-Based CLI

By default, two roles exist in all switches:

- Network operator—Has permission to view the configuration.
- Network administrator—Has permission to execute all commands and to set up to 64 permission levels based on user roles and groups (see [Chapter 16, “Configuring Switch Security”](#)).

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have the correct permission as specified in the description of the command.

Using Valid Formats and Ranges


Note

Do not enter ellipsis (...), vertical bar (|), less or great (<>), bracket ([]), or braces ({ }) in command lines. These characters have special meaning in SAN-OS text strings.

Some commands require a MAC address, IP address, or IDs that must be designated in a standard format or given a range. See [Table 2-4](#).

Table 2-4 Valid Formats and Ranges

Address	Description	Valid Format Example	Range
MAC address	6 bytes in hexadecimal format separated by colons (not case-sensitive)	00:00:0c:24:d2:Fe	—
IP address	32 bytes, written as 4 octets separated by periods (dotted decimal format) that are made up of a network section, an optional netmask section, and a host section.	126.2.54.1	—
VSAN	Integer that specifies the VSAN.	7	1 to 4093
VLAN	Integer that specifies the VLAN	11	1 to 4093
Port WWN (pWWN)	Eight hexadecimal numbers separated by colons (not case-sensitive).	12:34:56:78:9A:BC:dE:F1	—
Node WWN (nWWN)	Eight hexadecimal numbers separated by colons (not case-sensitive).	12:34:56:78:9A:BC:dE:F1	—
LUN	8 bytes in hexadecimal format separated by colons. A minimum of two hex characters are acceptable. The valid format is hhhh[:hhhh[:hhhh[:hhhh]]]	64 (100d = 64h)	—
FC ID	Six character hexadecimal value prepended by 0x.	0xabc123	—
Domain ID	Integer that specifies the domain.	7	1 to 239
Timers	Integer that specifies timers in milliseconds for latency, FC time out values (TOV).	100	0 to 2147483647
Switching module	Slot in which the applicable switching module resides.	1	1 to 15
Switch priority	Integer specifying switch priority.	5	1 to 254
Channel group	Integer that specifies a PortChannel group addition.	1	1 to 100
Fabric Shortest Path First (FSPF)	Integer that specifies the hold time (in milliseconds) before making FSPF computations.	1000	0 to 65535
Fabric Analyzer	The allowed range for the frame size limit in bytes.	64	64 to 65536
Fabric Analyzer captures	An example of 10 frames, limits the number of frames captured to 10.	10	0 to 2147483647
FCIP profile	Integer that specifies the FCIP profile	101	1 to 255
TCP retransmit time	Integer that specifies the minimum retransmit time for the TCP connection in milliseconds	300	250 to 5000

Table 2-4 Valid Formats and Ranges (continued)

Address	Description	Valid Format Example	Range
Keepalive timeout	Integer that specifies the TCP connection's keepalive timeout in seconds.	60	1 to 7200
TCP retransmissions	Integer that specifies the maximum number of TCP transmissions.	6	1 to 8
PMTU	Integer that specifies the path MTU reset time in seconds	90	60 to 3600
TCP buffer size	Integer that specifies the advertised TCP buffer size in KB.	5000	0 to 8192
Traffic burst size	Integer that specifies the maximum burst size in KB.	30	10 to 100
Peer TCP port	Integer that specifies the TCP port number	3000	0 to 65535
Acceptable time difference	Integer that specifies the acceptable time difference in milliseconds for a packet being accepted.	4000	1 to 60,000
iSCSI pWWN allocation	Integer that specifies the number of pWWNs that must be allocated to an iSCSI initiator.	2	1 to 64
CDP refresh and hold time	Integer that specifies the refresh time interval and the hold time in seconds for the CDP protocol.	60	5 to 255



Obtaining and Installing Licenses

The licensing functionality is available in all switches in the Cisco MDS 9000 Family. This functionality allows you to access specified premium features on the switch after you install the appropriate license for that feature. Licenses are sold, supported, and enforced from Release 1.3(1).

This section contains information related to licensing types, licensing procedure, license installation, and license management for the Cisco MDS SAN-OS software.

This chapter includes the following sections:

- [License Terminology, page 3-2](#)
- [Licensing Model, page 3-4](#)
- [Licensing High Availability, page 3-5](#)
- [Options to Install a License, page 3-5](#)
- [Obtaining a Factory-Installed License, page 3-6](#)
- [Performing a Manual Installation, page 3-6](#)
- [Obtaining License Key Files, page 3-7](#)
- [Installing the License Key File, page 3-7](#)
- [Uninstalling Licenses, page 3-8](#)
- [Updating Licenses, page 3-9](#)
- [License Expiry Alerts, page 3-10](#)
- [Moving Licenses Between Switches, page 3-10](#)
- [Displaying License Information, page 3-11](#)

License Terminology

The terms used in this chapter are explained in this section.

- **Licensed feature**
Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- **License expiry**
The time span during which a licensed feature is valid. The software tracks all licenses and sends periodic alerts before shutting down the licensed feature.
- **Counted license**
The number of usage instances for a licensed feature.
- **Licensed application**
A software feature that requires a license in order to be used.
- **License enforcement**
Mechanism to prevent features being used without first obtaining a license.
- **Node-locked license**
A license that can only be used on a particular switch using the switch's unique host ID.
- **Host ID**
A host ID is a unique chassis serial number that is specific to each Cisco MDS switch.
- **Proof of purchase**
Also known as the Claim Certificate. A document entitling its rightful owner to use licensed feature(s) of the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches on one MDS switch as described in that document.
- **Product Authorization Key (PAK)**
Using the PAK, you can obtain a license key from one of the sites listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions via e-mail.
- **License key file**
A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
 - License keys are required if your switch is running SAN-OS Release 1.3(x) or greater.
 - License keys are not required to use licensed features in SAN-OS software Release 1.(x) and 1.2(x).
- **Counted licenses**
Counted licenses refer to the number of licenses issued for a single feature (for example, FCIP). You can increase counted licenses (incremental licenses) should a need arise in the future.
- **Incremental licenses**
Incremental refers to adding other licensed features not in the initial license file. License keys are incremental—if you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.

- Evaluation license

An evaluation license is a temporary license. Evaluation licenses are time bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).

- Permanent license

A license that is not time bound (does not have an expiry date) is called a permanent license.

- Grace period

The grace period for a license refer to the amount of time an application can continue functioning without a license. The grace period is set to 60 days from the first occurrence of using any licensed feature without a license. See the “[License Expiry Alerts](#)” section on page 3-10 for more information. The grace period starts with the first checkout and will be counted regardless of the feature being used.

- Support

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Licensing Model

The licensing model defined for the Cisco MDS product line has two options:

- Feature-based licensing: features that are applicable to the entire switch. The cost varies based on a per-switch usage. [Table 3-1](#) lists the feature-based license packages.
- Module-based licensing: features that require additional hardware modules. The cost varies based on a per-module usage. An example is the IPS-8 module using the FCIP feature.

Table 3-1 Feature-Based Licenses

Feature License	Features
Standard package (free—no license required)	<ul style="list-style-type: none"> • FCP, SSH, SFTP, and iSCSI protocols • Fabric Manager and Remote monitoring (RMON) • VSANs, High availability, PortChannel, and Zoning • Fibre Channel Congestion Control (FCC) • Virtual Output Queuing (VOQ) • Diagnostics (SPAN, RSPAN, and FC Analyzer) • SNMP v3, Role-based access control, RADIUS • Call Home and Interoperability modes • IP access control lists (ACLs) • Terminal Access Controller Access Control System (TACACS+) • Fabric-Device Management Interface (FDMI) • Internet Storage Name Service (iSNS) client.
Enterprise package (ENTERPRISE_PKG)	<ul style="list-style-type: none"> • Enhanced security features: <ul style="list-style-type: none"> – LUN zoning – Read-only zones – Port security – VSAN-based access control – Fibre Channel Security Protocol (FC-SP) authentication • Advanced traffic engineering—Quality of Service (QoS) • Enhanced VSAN routing—inter-VSAN routing
SAN extension over IP (SAN_EXTN_OVER_IP)	<ul style="list-style-type: none"> • FCIP protocol for the IPS-8 module • FCIP compression for the IPS-8 module • FCIP write acceleration for the IPS-8 module

Table 3-1 Feature-Based Licenses (continued)

Feature License	Features
Mainframe (MAINFRAME_PKG)	<ul style="list-style-type: none"> • FICON protocol and CUP management • FICON VSAN and intermixing • Switch cascading • Fabric Binding
Fabric Manager Server (FM_SERVER_PKG)	<ul style="list-style-type: none"> • Multiple physical fabric management • Centralized fabric discovery services • Continuous MDS health and event monitoring • Long term historical Fibre Channel Performance monitoring • Performance reports and charting for hotspot analysis

Licensing High Availability

Similar to any other Cisco MDS SAN-OS feature, the licensing feature also maintains the following high availability standards for all switches in the Cisco MDS 9000 Family:

- Installing any license in any switch is a nondisruptive process.
- Installing a license automatically saves a copy of permanent licenses to the chassis in all switches.
- When a licensed feature is enabled without a license key, the MDS switch enables the feature and starts a counter on the grace period. You then have 60 days to install the appropriate license keys or disable the use of that feature. If at the end of the 60 day grace period the switch does not have a valid license key for the feature, the feature is automatically disabled by the switch.
- Directors in the Cisco MDS 9500 Series have the following additional high availability features
 - The license software runs on both supervisor modules and provides failover protection.
 - The license key file is mirrored on both supervisor modules. Even if both supervisor modules fail, the license file continues to function from the version that is available on the chassis.

Options to Install a License

If you have purchased a new switch through either your reseller or through Cisco, you have two options:

- To have the licenses preinstalled in the factory (see the [“Obtaining a Factory-Installed License” section on page 3-6](#)).
- To install the licenses yourself by following the manual process (see the [“Performing a Manual Installation” section on page 3-6](#)).

If you already have an existing switch, follow the manual process (see the [“Performing a Manual Installation” section on page 3-6](#)).

Obtaining a Factory-Installed License

You can obtain factory-installed licenses for a new switch.

To obtain a factory-installed license for a new Cisco MDS switch, follow these steps.

Step 1 Contact your reseller or Cisco representative and request this service.



Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Your switch is shipped with the required licenses installed in the system. The Proof of Purchase document is sent along with the switch.

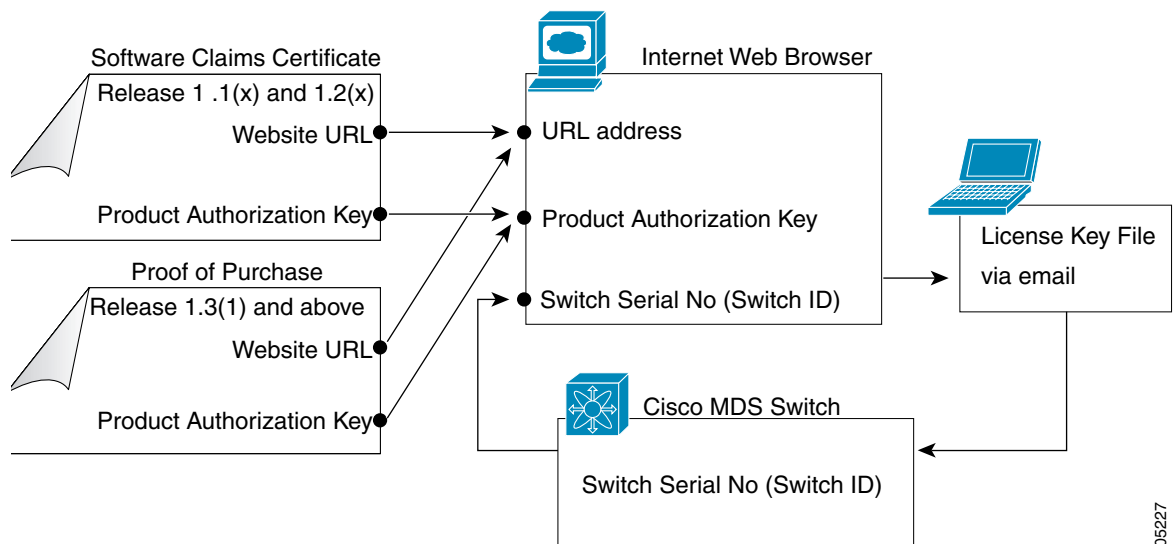
Step 2 Obtain the host ID from the Proof of Purchase for future use.

Step 3 Start using the switch and the installed licenses features.

Performing a Manual Installation

If you have existing switches or if you wish to install the licenses on your own, you must first obtain the license key file (see the “[Obtaining License Key Files](#)” section on page 3-7) and then install that file (see the “[Installing the License Key File](#)” section on page 3-7) in the switch (see [Figure 3-1](#)).

Figure 3-1 Obtaining a License Key File



105227

Obtaining License Key Files

To obtain new or updated license key files, follow these steps.

-
- Step 1** Use the **show license host-id** command to obtain the serial number for your switch.
The host ID is also referred to as the switch serial number.
- Step 2** Obtain your Claim Certificate or the Proof of Purchase document.
This document accompanies every Cisco MDS switch.
- Step 3** Locate the Product Authorization Key (PAK) from the Claim Certificate or Proof of Purchase document.
- Step 4** Locate the website URL from the Claim Certificate or Proof of Purchase document.
- Step 5** Access the specified URL that applies to your switch and enter the switch serial number and the PAK.
The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the switch for which it was requested. The requested features are also enabled once the SAN-OS software on the specified switch access the license key file.

**Caution**

Install the license file in the specified MDS switch without making any modifications.

A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for using that license starts from the first time you start using a feature offered by that license (see the [“License Expiry Alerts”](#) section on page 3-10 for further information).

**Note**

Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for details on installing automated licenses using the Fabric Manager GUI.

Installing the License Key File

**Tip**

If you need to install multiple licenses in any switch in the Cisco MDS 9000 Family, be sure to provide unique file names for each license key file.

To install a license key file in any switch, follow these steps:

-
- Step 1** Log into the switch through the console port of the active supervisor.
- Step 2** Perform the installation by issuing the **install license** command on the active supervisor module from the switch console.

```
switch# install license bootflash:license_file.lic
Installing license ..done
```

**Note**

If you provide a target name for the license key file, the file is installed with the specified name. Otherwise, the file name specified in the license key file is used to install the license.

- Step 3** Exit the switch console and open a new terminal session to view all license files installed on the switch using the **show license** command.

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
    HOSTID=VDH=FOX0646S017 \
    NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
    <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

**Note**

If the license meets all guidelines when the **install license** command is issued, all features and modules continue functioning as configured. This is true for any switch in the Cisco MDS 9000 Family.

Uninstalling Licenses

You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is currently being used, the software rejects the request and issues an error message. Uninstalling an unused license causes the grace period to come into effect. The grace period is counted from the first use of the feature without a license and is reset when a valid license file is installed.

**Note**

Permanent licenses cannot be uninstalled if they are currently being used. Features turned on by permanent licenses must first be disabled, before that license is uninstalled.

**Tip**

If you are using an evaluation license and would like to install a new permanent license, you can do so without service disruption and before the evaluation license expires. Removing an evaluation license immediately triggers a grace period without service disruption.

**Caution**

Uninstalling a license requires the related features to first be disabled.

To uninstall a license, follow these steps:

- Step 1** Save your running configuration to a remote server using the **copy** command (see the [“Copying Files” section on page 4-27](#)).
- Step 2** Disable the features provided by the license to be uninstalled. Issue the **show license usage package_name** command to view the enabled features for a specified package.
- ```
switch# show license usage ENTERPRISE_PKG
Application

ivr
qos_manager

```
- Step 3** Issue the **show license brief** command in EXEC mode to view a list of all installed license files and identify the file to be uninstalled. In this example, the file to be uninstalled is the Ficon.lic file.

```
switch# show license brief
Enterprise.lic
Ficon.lic
```

- Step 4** Uninstall the Ficon.lic file using the **clear license filename** command, where *filename* is the name of the installed license file.

```
switch# clear license Ficon.lic
Clearing license Ficon.lic:
SERVER this_host ANY
VENDOR cisco
An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
 NOTICE=<LicFileID>san_extn2.lic</LicFileID><LicLineID>1</LicLineID> \
 SIGN=67CB2A8CCAC2
```

- Step 5** Enter **yes** (yes is the default), to continue with the license update.

```
Do you want to continue? (y/n) y
Clearing license ..done
switch#
The Ficon.lic license file is now uninstalled.
```

## Updating Licenses

If your license is time bound, you must obtain and install an updated license. Contact technical support to request an updated license.



### Note

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:  
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

To update a license, follow these steps:

- Step 1** Obtain the updated license file using the procedure described in the “Obtaining License Key Files” section on page 3-7.
- Step 2** Save your running configuration to a remote server using the **copy** command (see the “Copying Files” section on page 4-27).
- Step 3** Issue the **show license brief** command to verify the name of the file to be updated.
- Step 4** Update the license file using the **update license url** command, where *url* specifies the bootflash:, slot0:, or volatile: directory location of the updated license file.

```
switch# show license brief
sanextn1.lic:

switch# update license bootflash:sanextn2.lic sanextn1.lic
Updating sanextn1.lic:
SERVER this_host ANY
VENDOR cisco
An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
 NOTICE=<LicFileID>san_extn1.lic</LicFileID><LicLineID>0</LicLineID> \
 SIGN=33088E76F668
```

```

with bootflash:/sanextn2.lic:
SERVER this_host ANY
VENDOR cisco
An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
 NOTICE=<LicFileID>san_extn2.lic</LicFileID><LicLineID>1</LicLineID> \
 SIGN=67CB2A8CCAC2

```

**Step 5** Enter **yes** (yes is the default), to continue with the license update.

```

Do you want to continue? (y/n) y
Updating license ..done
switch#
The sanextn1.lic license file is now updated.

```

## License Expiry Alerts

The SAN-OS license counter keeps track of all licenses on a switch. Once an expiry period has started, you will receive CLI console messages, SNMP traps, syslog error messages, and Call Home messages on a daily basis.

Beyond that, the frequency of these message will increase to an hourly basis during the last seven days of the expiry time span. For example:

You FICON license feature is scheduled to expire in 60 days. If today is December 1st, the license expires on January 30th. In this case, you will receive:

- Daily alerts from December 1st to January 23rd
- Hourly alerts from January 24th to January 29th
- From January 30th, the FICON feature will run without a license for a grace period of 60 days.
- From January 30th to March 21st, you will receive daily alerts about the grace period usage.
- From March 22nd to March 30th, you will receive hourly alerts about the grace period ending.
- On March 31st, the FICON feature is automatically turned off.



### Note

License expiry alerts cannot be configured.



### Caution

After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. The grace period also applies to licensed features in Release 1.2(x). While Release 1.2(x) did not enforce the licenses, any upgrade will enforce license requirements and the 60-day grace period.

## Moving Licenses Between Switches

A license is specific to the switch for which it is issued and is not valid on any other switch. If you need to transfer a license from one switch to another, contact your customer service representative.

**Note**

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:  
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

## Displaying License Information

Use the **show license** commands to display all license information configured on this switch (see Examples 3-4 to 3-6).

### Example 3-1 Displays Information About Current License Usage

```
switch# show license usage
Feature Installed License Status ExpiryDate Comments
 Count

FM_SERVER_PKG Yes - Unused never license missing
MAINFRAME_PKG No - Unused
ENTERPRISE_PKG Yes - InUse never -
SAN_EXTN_OVER_IP No 0 Unused
**** WARNING: License file(s) missing. ****
```

### Example 3-2 Displays the List of Features in a Specified Package

```
switch# show license usage ENTERPRISE_PKG
Application

ivr
qos_manager

```

### Example 3-3 Displays the Host ID for the License

```
switch# show license host-id
License hostid: 1VDH=FOX0646S017
```

### Example 3-4 Displays All Installed License Files and their Contents

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
 HOSTID=VDH=FOX0646S017 \
 NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
 <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
Evaluation.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 30-Dec-2003 uncounted \
 HOSTID=VDH=FOX0646S017 \
 NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
 <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

1. VDH = Vendor Defined Host-ID

**Example 3-5 Displays a List of Installed License File Names**

```
switch# show license brief
Enterprise.lic
Ficon.lic
FCIP.lic
```

**Example 3-6 Displays the Contents of a Specified License File**

```
switch# show license file Permanent.lic
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
 HOSTID=VDH=FOX0646S017 \
 NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
 <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```





## Initial Configuration

---

This chapter describes how to initially configure switches so they can be accessed by other devices. This chapter includes the following sections:

- [Starting a Switch in the Cisco MDS 9000 Family, page 4-2](#)
- [Initial Setup Routine, page 4-2](#)
- [Assigning a Switch Name, page 4-13](#)
- [Accessing the Switch, page 4-14](#)
- [Where Do You Go Next?, page 4-14](#)
- [Verifying the Module Status, page 4-15](#)
- [Configuring Date and Time, page 4-15](#)
- [Configuring the Management Port, page 4-20](#)
- [Disabling a Telnet Server, page 4-22](#)
- [Working with Configuration Files, page 4-23](#)
- [Copying Files, page 4-27](#)
- [Deleting Files, page 4-30](#)
- [Configuring Console Settings, page 4-31](#)
- [Configuring COM1 Settings, page 4-32](#)
- [Configuring Modem Connections, page 4-33](#)
- [Configuring CDP, page 4-37](#)

# Starting a Switch in the Cisco MDS 9000 Family

The following procedure is a review of the tasks you should have completed during hardware installation, including starting up the switch. These tasks must be completed before you can configure the switch.

- 
- Step 1** Check that the switch is set for the correct AC (or DC) power voltages.  
Refer to either the *Cisco MDS 9000 Family Hardware Installation Guide* (for the required product) for correct power voltages.
- Step 2** Connect the power cord(s) to the switch.
- Step 3** Connect the console port to the switch.



---

**Note** The console port is an asynchronous (async) serial port; any device connected to this port must be capable of asynchronous transmission.

---

Before connecting the console port, check the terminal documentation to determine the baud rate. The baud rate of the terminal must match the default baud rate (9600 baud) of the console port. Set up the terminal as follows (see the [“Configuring Console Settings” section on page 4-31](#)):

- 9600 baud
- 8 data bits
- 1 stop bit
- No parity

- Step 4** Power on the switch. The switch boots automatically.
- 

## Initial Setup Routine

The first time that you access a switch in the Cisco MDS 9000 Family, it runs a setup program that prompts you for the IP address and other configuration information necessary for the switch to communicate over the supervisor module Ethernet interface. This information is also required if you plan to configure and manage the switch.



---

**Note** The IP address must first be set up in CLI when the switch is powered up for the first time so the Cisco MDS 9000 Fabric Manager can reach the switch.

---

## Preparing to Configure the Switch

Before you configure a switch in the Cisco MDS 9000 Family for the first time, you need the following information:

- Administrator password—you have the following options:
  - Change the default password (admin) for the administrator.
  - Create an additional login account and password.
- SNMPv3 user name and authentication password.
- SNMP community string.
- Switch name—this is your switch prompt.
- IP address for the switch's management interface—The management interface can be an out-of-band Ethernet interface or an in-band Fibre Channel interface.
- Subnet mask for the switch's management interface.
- The following IP addresses:
  - Destination prefix, destination prefix subnet mask, and next hop IP address, if you want to enable IP routing. Also, provide the IP address of the default network.
  - Otherwise, provide an IP address of the default gateway.
- DNS IP address (optional).
- Default domain name (optional).
- SSH service on the switch—if you wish to enable this service, then select the type of SSH key (dsa/rsa/rsa1) and number of key bits (768 to 2048).
- NTP server IP address (optional).

**Note**

Be sure to configure the IP routing, the IP default network address, and the IP default gateway address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

## Default Login

All Cisco MDS 9000 family switches have the network administrator as a default user (admin) and a default password (admin). You can change the default password, if required, during the initial setup process. You cannot change the default user at any time.

During the initial setup process, you have the option to configure one additional user in the network administrator role. See the [“Configuring Role-Based CLI Authorization”](#) section on page 16-18 for information of default roles and permissions.

If you change the administrator password during the initial setup process and subsequently forget this new password, you have the option to recover this password (see the [“Recovering Administrator Password”](#) section on page 16-26).

## Setup Options

The setup scenario differs based on the subnet to which you are adding the new switch. You must configure a switch in the Cisco MDS 9000 Family with an IP address to enable management connections from outside of the switch.

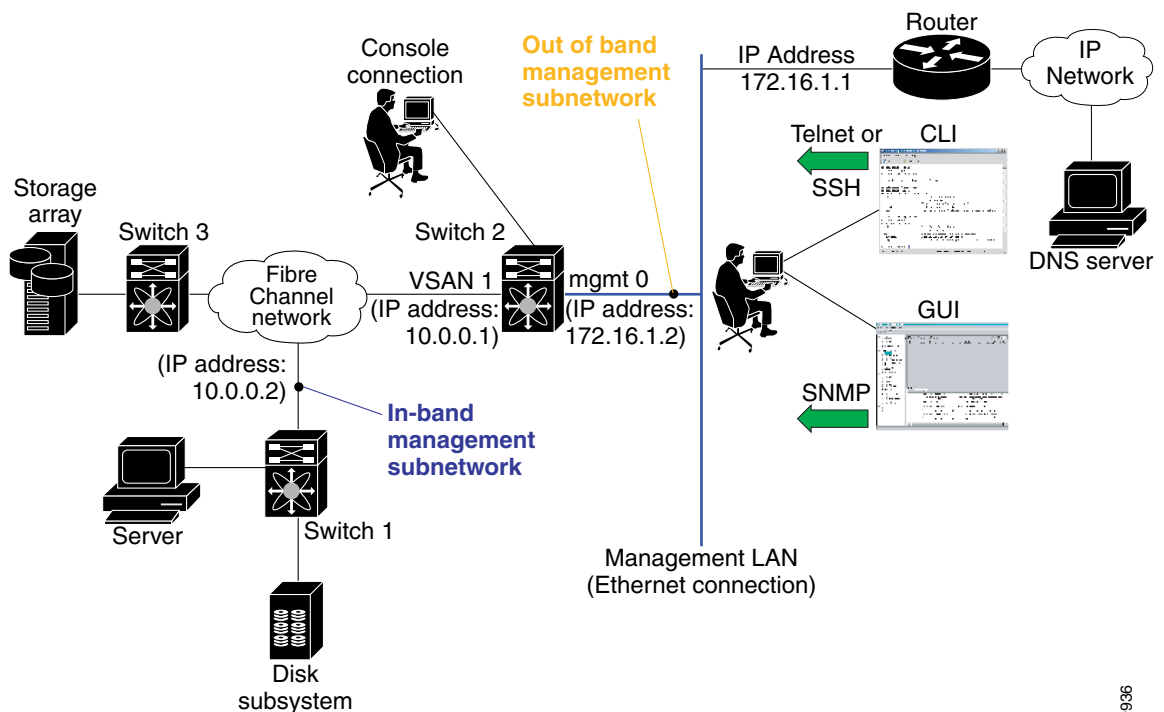


### Note

Some concepts like out-of-band management and in-band management are briefly explained here. These concepts are explained in more detail in subsequent chapters.

- Out-of-band management—This feature provides a connection to the network through a supervisor module front panel Ethernet port (see [Figure 4-1](#)).
- In-band management—This feature provides IP over Fibre Channel (IPFC) to manage the switches. The in-band management feature is transparent to the network management system (NMS). Instead of conventional Ethernet physical media, switches in the Cisco MDS 9000 Family use IPFC as the transport mechanism (see [Figure 4-1](#) and [Chapter 20, “Configuring IP Services”](#)).

**Figure 4-1 Management Access to Switches**



79936

## Assigning Setup Information

This section describes how to initially configure the switch for both out-of-band and in-band management.



### Note

Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point.



### Tip

If you do not wish to answer a previously-configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example switch name), the switch uses what is previously configured and skips to the next question.

## Configuring Out-of-Band Management



### Note

If you wish to configure both in-band and out-of-band configuration together, enter **Yes** in both Step 11 and Step 12 in the following procedure.

To configure the switch for first time out-of-band access, follow these steps:

**Step 1** Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

**Step 2** Enter **yes** to enter the setup mode.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt, to end the configuration process.

**Step 3** Enter the new password for the administrator (admin is the default):

Enter the password for admin: **admin**

**Step 4** Enter **yes** (no is the default), to create additional accounts.

Create another login account (yes/no) [n]: **yes**

While configuring your initial setup, you can create an additional user account (in the network-admin role) besides the administrator's account. See the [“Configuring Role-Based CLI Authorization”](#) section on page 16-18 for information of default roles and permissions.

a. Enter the user login ID.

Enter the user login ID: *user\_name*

b. Enter the user password.

Enter the password for user\_name: *user-password*

**Step 5** Enter **yes** (yes is the default), if you wish to create an SNMPv3 account.

Configure SNMPv3 Management parameters (yes/no) [y]: **yes**

**a.** Enter the user name (admin is the default).

SNMPv3 user name [admin]: **admin**

**b.** Enter the SNMPv3 password (minimum of eight characters).

SNMPv3 user authentication password : **admin\_pass**



**Note**

If this password contains less than eight characters, you must enter a new password. The same password is also used for the SNMPv3 privacy.

By default, if the admin password is at least eight characters, then the SNMP authentication password will be same as admin password (at least eight characters). If the admin password is less than eight characters, then you need to provide a new password for SNMP.

The admin password can have a minimum of one character, but the SNMP authentication password must have a minimum of eight characters.

**Step 6** Enter **yes** (no is the default) to configure read-only or read-write SNMP community string.

Configure read-only SNMP community string (yes/no) [n]: **yes**

**a.** Enter the SNMP community string.

SNMP community string: **snmp\_community**

**Step 7** Enter a name for the switch.



**Note**

The switch name is limited to 32 alphanumeric characters.

Enter the switch name: **switch\_name**

**Step 8** Enter **yes** (yes is the default) to configure out-of-band management.

Continue with Out-of-band (mgmt0) management configuration? [yes/no]: **yes**

**a.** Enter the mgmt0 IP address.

Mgmt0 IP address: **ip\_address**

**b.** Enter the mgmt0 subnet mask.

Mgmt0 IP netmask: **subnet\_mask**

**Step 9** Enter **no** (no is the default) at the in-band management configuration prompt.

Continue with in-band (VSAN1) management configuration? (yes/no) [no]: **no**

**Step 10** Enter **yes** (yes is the default) to enable IP routing capabilities.

Enable the ip routing capabilities? (yes/no) [y]: **yes**

**Step 11** Enter **yes** (yes is the default) to configure a static route (recommended).

Configure static route: (yes/no) [y]: **yes**

**a.** Enter the destination prefix.

Destination prefix: **dest\_prefix**

- b. Type the destination prefix mask.

Destination prefix mask: *dest\_mask*

- c. Type the next hop ip address.

Next hop ip address: *next\_hop\_address*



**Note**

Be sure to configure the IP routing, the IP default network address, and the IP default gateway address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

- Step 12** Enter **yes** (yes is the default) to configure the default-network (recommended).

Configure the default-network: (yes/no) [y]: **yes**

- a. Enter the default-network IP address.



**Note**

The default network address is the destination prefix provided in Step 10 a above.

Default network IP address: *dest\_prefix*

- Step 13** Enter **yes** (yes is the default) to configure the default-gateway (recommended).

Configure the default-gateway: (yes/no) [y]: **yes**

- a. Enter the default-gateway IP address.

IP address of the default-gateway: *default\_gateway*

- Step 14** Enter **yes** (yes is the default) to configure the DNS IP address.

Configure the DNS IP address? (yes/no) [y]: **yes**

- a. Enter the DNS IP address.

DNS IP address: *name\_server*

- Step 15** Enter **yes** (default is no) to configure the default domain name.

Configure the default domain name? (yes/no) [n]: **yes**

- a. Enter the default domain name.

Default domain name: *domain\_name*

- Step 16** Enter **yes** (yes is the default), to enable Telnet service.

Enable the telnet service? (yes/no) [y]: **yes**

- Step 17** Enter **yes** (no is the default) to enable the SSH service.

Enabled SSH service? (yes/no) [n]: **yes**

- Step 18** Enter the SSH key type (see [Generating an SSH Host Key Pair, page 16-27](#)) you would like to generate.

Type the SSH key you would like to generate (dsa/rsa/rsa1)? **dsa**

- Step 19** Enter the number of key bits within the specified range.

Enter the number of key bits? (512 to 2048): **768**

**Step 20** Enter **yes** (no is the default) to configure the NTP server.

Configure NTP server? (yes/no) [n]: **yes**

a. Enter the NTP server IP address.

NTP server IP address: *ntp\_server\_IP\_address*

**Step 21** Enter **shut** (shut is the default) to configure the default switchport interface to the shut state.

Configure default switchport interface state (shut/noshut) [shut]: **shut**

**Step 22** Enter **on** (on is the default) to configure the switchport trunk mode.

Configure default switchport trunk mode (on/off/auto) [on]: **on**

**Step 23** Enter **permit** (deny is the default) to permit a default zone policy.

Configure default zone policy (permit/deny) [deny]: **permit**

Permits traffic to flow to all members of the default zone.

**Step 24** Enter **yes** (no is the default) to enable a full zoneset distribution (see the [“Distributing Zone Sets”](#) section on page 13-11).

Enable full zoneset distribution (yes/no) [n]: **yes**

Enables full-zoneset distribution throughout the fabric.

**Step 25** Enter **yes** (no is the default) to enable FC ID persistence (see the [“Enabling Persistent FC IDs”](#) section on page 24-9).

Enable FCID Persistence globally (yes/no) [n]: **yes**

Enables FC ID persistence throughout the fabric.

**Step 26** Review and edit the configuration that you have just entered.

**Step 27** Enter **no** (no is the default) if you are satisfied with the configuration.

The following configuration will be applied:

```
username admin password admin_pass role network-admin
username user_name password user_pass role network-admin
snmp-server user admin network-admin auth md5 admin_pass priv admin_pass
snmp-server community snmp_community ro
switchname switch
interface mgmt0
 ip address ip_address subnet_mask
 no shutdown
ip routing
ip route dest_prefix dest_mask dest_address
ip default-network dest_prefix
ip default-gateway default_gateway
ip name-server name_server
ip domain-name domain_name
telnet server enable
ssh key dsa 768 force
ssh server enable
ntp server ipaddr ntp_server
system default switchport shutdown
system default switchport trunk mode on
zone default-zone permit vsan 1-4093
zoneset distribute full vsan 1-4093
fcdomain fcid persistent global-enable
fcdomain fcid persistent vsan 1
```



Would you like to edit the configuration? (yes/no) [n]:

**Step 28** Enter **yes** (yes is default) to use and save this configuration:

Use this configuration and save it? (yes/no) [y]: **yes**



#### Caution

If you do not save the configuration at this point, none of your changes will be updated the next time the switch is rebooted. Ensure to type **yes** in order to save the new configuration. This will ensure that the kickstart and system boot images are also automatically configured (see [Chapter 6, “Software Images”](#)).

## In-Band Management Configuration

The in-band management logical interface is VSAN 1. This management interface uses the Fibre Channel infrastructure to transport IP traffic. An interface for VSAN 1 is created on every switch in the fabric. Each switch should have its VSAN 1 interface configured with an IP address in the same subnetwork. A default route, pointing to the switch that provides access to the IP network should be configured on every switch in the Fibre Channel fabric (see [Chapter 9, “Configuring and Managing VSANs”](#)).



#### Note

If you wish to configure both in-band and out-of-band configuration together, enter **Yes** in both Step 11 and Step 12 in the following procedure.

To configure a switch for first time in-band access, follow these steps:

**Step 1** Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

**Step 2** Enter **yes** to enter the setup mode.

```
---- Basic System Configuration Dialog ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process. Press **Ctrl-C** from any prompt, to abort the configuration process.

**Step 3** Enter the new password for the administrator.

Enter the password for admin: admin

**Step 4** Enter **no** (no is the default), if you do not wish to create other additional accounts.

Create another login account (yes/no) [no]: **no**

**Step 5** Enter **yes** (yes is the default), if you wish to create a SNMPv3 account.

Configure SNMPv3 Management parameters (yes/no) [y]: **yes**

- a. Enter the user name.

```
SNMPv3 user name [admin]: user_name
```

By default, the SNMP user name is admin.

- b. Enter the SNMPv3 password (minimum of 8 characters).

```
SNMPv3 user authentication password [admin_pass]: admin_pass
```



**Note**

If this password contains less than eight characters, you must enter a new password. The same password is also used for the SNMPv3 privacy.

By default, if the admin password is at least eight characters, then the SNMP authentication password will be same as admin password (at least eight characters). If the admin password is less than eight characters, then you need to provide a new password for SNMP.

The admin password can have a minimum of one character, but the SNMP authentication password must have a minimum of eight characters.

**Step 6** Configure read-only or read-write SNMP community string.

- a. Enter **no** (no is the default) to avoid configuring read-only SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: no
```

- b. Enter **no** (no is the default) to configure read-only SNMP community string.

```
Configure read-only SNMP community string (yes/no) [n]: yes
```

- c. Enter the SNMP community string.

```
SNMP community string: snmp_community
```

**Step 7** Enter a name for the switch.



**Note**

The switch name is limited to 32 alphanumeric characters.

```
Enter the switch name: switch_name
```

**Step 8** Enter **no** (yes is the default) at the configuration prompt to configure out-of-band management.

```
Continue with Out-of-band (mgmt0) management configuration? [yes/no]: no
```

**Step 9** Enter **yes** (no is the default) at the in-band management configuration prompt.

```
Continue with in-band (VSAN1) management configuration? (yes/no) [no]: yes
```

- a. Enter the VSAN 1 IP address.

```
VSAN1 IP address: ip_address
```

- b. Enter the subnet mask.

```
VSAN1 IP net mask: subnet_mask
```

**Step 10** Enter **yes** (yes is the default) to enable IP routing capabilities.

```
Enable ip routing capabilities? (yes/no) [y]: yes
```

**Step 11** Enter **no** (yes is the default) to configure a static route.

```
Configure static route: (yes/no) [y]: no
```

- Step 12** Enter **yes** (yes is the default) to configure the default network.  
Configure the default-network: (yes/no) [y]: **no**
- Step 13** Enter **yes** (yes is the default) to configure the default gateway.  
Configure the default-gateway: (yes/no) [y]: **yes**
- a.** Enter the default-gateway IP address.  
IP address of the default-gateway: default\_gateway
- Step 14** Enter **no** (yes is the default) to configure the DNS IP address.  
Configure the DNS IP address? (yes/no) [y]: **no**
- Step 15** Enter **no** (no is the default) to skip the default domain name configuration.  
Configure the default domain name? (yes/no) [n]: **no**
- Step 16** Enter **no** (yes is the default), to disable Telnet service.  
Enable the telnet service? (yes/no) [y]: **no**
- Step 17** Enter **yes** (no is the default) to enable the SSH service.  
Enabled SSH service? (yes/no) [n]: **yes**
- Step 18** Enter the SSH key type (see [Generating an SSH Host Key Pair, page 16-27](#)) you would like to generate.  
Type the SSH key you would like to generate (dsa/rsa/rsal)? **rsa**
- Step 19** Enter the number of key bits within the specified range.  
Enter the number of key bits? (512 to 1024): **1024**
- Step 20** Enter **no** (no is the default) to configure the NTP server.  
Configure NTP server? (yes/no) [n]: **no**
- Step 21** Enter **noshut** (shut is the default) to configure the default switchport interface to the up state.  
Configure default switchport interface state (shut/noshut) [shut]: **noshut**
- Step 22** Enter **auto** (on is the default) to configure the switchport trunk mode automatically.  
Configure default switchport trunk mode (on/off/auto) [on]: **auto**
- Step 23** Enter **deny** (deny is the default) to deny a default zone policy.  
Configure default zone policy (permit/deny) [deny]: **deny**
- Step 24** Enter **no** (no is the default) to disable full zoneset distribution (see the [“Distributing Zone Sets” section on page 13-11](#)).  
Enable full zoneset distribution (yes/no) [n]: **no**
- Step 25** Enter **no** (no is the default) to enable FC ID persistence (see the [“Enabling Persistent FC IDs” section on page 24-9](#)).  
Enable FCID Persistence globally (yes/no) [n]: **no**
- The new configuration is displayed.
- Step 26** Review and edit the configuration that you have just entered.
- Step 27** Enter **no** (no is the default) if you are satisfied with the configuration.

The following configuration will be applied:

```
username admin password admin_pass role network-admin
snmp-server user snmp_user network-admin auth md5 snmp_pass priv snmp_pass
snmp-server community snmp_community rw
switchname switch
interface vsan1
 ip address ip_address subnet_mask
 no shutdown
ip default-gateway default_gateway
no telnet server enable
ssh key rsa 1024 force
ssh server enable
no system default switchport shutdown
system default switchport trunk mode auto
no zone default-zone permit vsan 1-4093
no zoneset distribute full vsan 1-4093
no fcdomain fcid persistent global-enable
```

Would you like to edit the configuration? (yes/no) [n]: **no**

**Step 28** Enter **yes** (yes is default) to use and save this configuration.

Use this configuration and save it? (yes/no) [y]: **yes**



**Caution**

If you do not save the configuration at this point, none of your changes will be updated the next time the switch is rebooted. Ensure to type **yes** in order to save the new configuration. This will ensure that the kickstart and system boot images are also automatically configured (see [Chapter 6, “Software Images”](#)).

## Using the setup Command

If you wish to make changes to the initial configuration at a later time, you can issue the **setup** command in EXEC mode.

```
switch# setup
---- Basic System Configuration Dialog ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

\*Note: setup always assumes a predefined defaults irrespective of the current system configuration when invoked from CLI.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process.

## Assigning a Switch Name

Each switch in the fabric requires a unique name. You can assign names to easily identify the switch by its physical location, its SAN association, or the organization to which it is deployed. The assigned name is displayed in the command-line prompt.



### Note

The switch name is limited to 32 alphanumeric characters.

This guide refers to a switch in the Cisco MDS 9000 Family as *switch*, and uses the `switch#` prompt.

To change the name of the switch, follow these steps:

|        | Command                                                                              | Purpose                                                                         |
|--------|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                        | Enters configuration mode.                                                      |
| Step 2 | <code>switch(config)# switchname myswitch1</code><br><code>myswitch1(config)#</code> | Changes the switch name prompt as specified.                                    |
| Step 3 | <code>myswitch1(config)# no switchname</code><br><code>switch(config)#</code>        | Reverts the switch name prompt to its factory default ( <code>switch#</code> ). |

## Assigning SNMP Switch Contact Information

Use the **snmp-server** command to set the contact information, switch location, and switch name. They are each limited to 32 characters (without spaces). Use the **no** form of the command to remove the system contact information. For more information on other **snmp-server** commands see the [“SNMP Security” section on page 16-30](#).

To configure contact information, follow these steps:

|        | Command                                                                                      | Purpose                                  |
|--------|----------------------------------------------------------------------------------------------|------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                                | Enters configuration mode.               |
| Step 2 | <code>switch(config)# snmp-server contact NewUser</code><br><code>switch(config)#</code>     | Assigns the contact name for the switch. |
|        | <code>switch(config)# no snmp-server contact NewUser</code><br><code>switch(config)#</code>  | Deletes the contact name for the switch. |
| Step 3 | <code>switch(config)# snmp-server location SanJose</code><br><code>switch(config)#</code>    | Assigns the switch location.             |
|        | <code>switch(config)# no snmp-server location SanJose</code><br><code>switch(config)#</code> | Deletes the switch location.             |

## Accessing the Switch

After initial configuration, you can access the switch in one of three ways (see [Figure 4-2](#)):

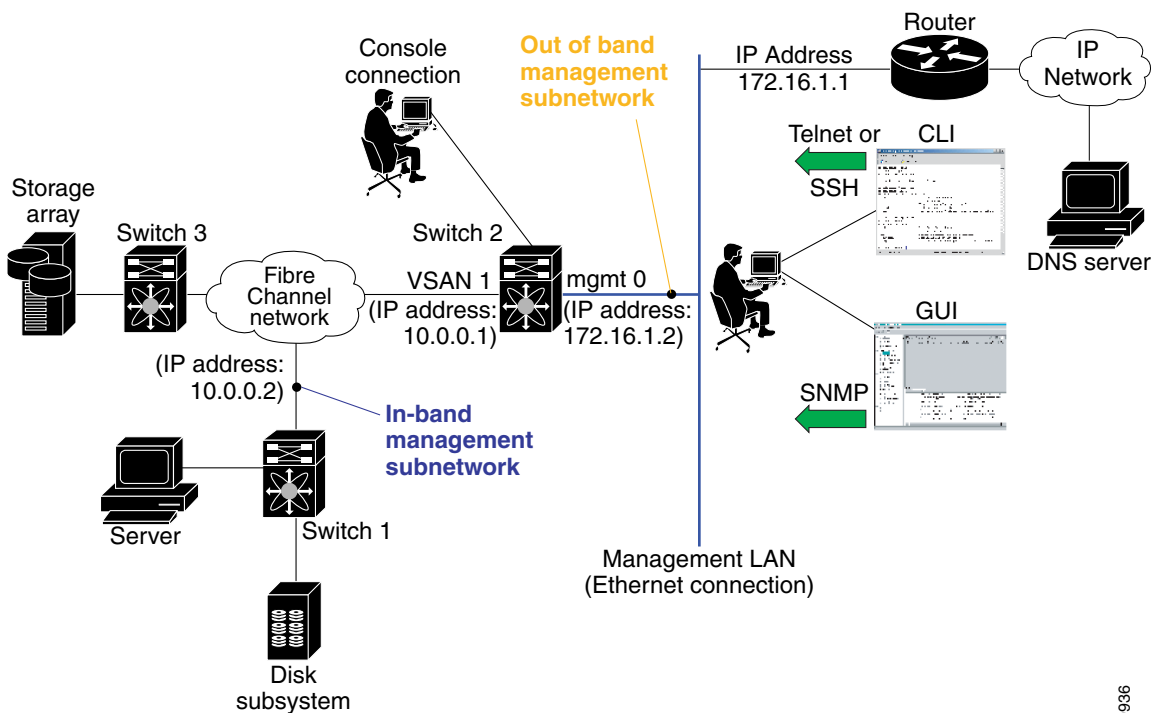
- In-band IP (IPFC) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use SNMP to connect to a Cisco MDS 9000 Fabric Manager GUI.
- Out-of-band (10/100BASE-T Ethernet) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use SNMP to connect to a Cisco MDS 9000 Fabric Manager GUI.


**Note**

To use the Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

- Serial console access—You can use a serial port connection to access the CLI.

**Figure 4-2 Switch Access Options**



79936

## Where Do You Go Next?

After reviewing the default configuration, you can change it or perform other configuration or management tasks. The initial setup can only be performed at the CLI. However, you can continue to configure other software features, or access the switch after initial configuration by using either the CLI or the Element Manager and Fabric Manager GUIs.

To use the Cisco MDS 9000 Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

## Verifying the Module Status

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
```

| Mod | Ports | Module-Type                | Model           | Status     |
|-----|-------|----------------------------|-----------------|------------|
| 2   | 8     | IP Storage Services Module | DS-X9308-SMIP   | ok         |
| 5   | 0     | Supervisor/Fabric-1        | DS-X9530-SF1-K9 | active *   |
| 6   | 0     | Supervisor/Fabric-1        | DS-X9530-SF1-K9 | ha-standby |
| 8   | 0     | Caching Services Module    | DS-X9560-SMAP   | ok         |
| 9   | 32    | 1/2 Gbps FC Module         | DS-X9032        | ok         |

| Mod | Sw          | Hw    | World-Wide-Name(s) (WWN)                           |
|-----|-------------|-------|----------------------------------------------------|
| 2   | 1.3(0.106a) | 0.206 | 20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00 |
| 5   | 1.3(0.106a) | 0.602 | --                                                 |
| 6   | 1.3(0.106a) | 0.602 | --                                                 |
| 8   | 1.3(0.106a) | 0.702 | --                                                 |
| 9   | 1.3(0.106a) | 0.3   | 22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00 |

| Mod | MAC-Address(es)                        | Serial-Num  |
|-----|----------------------------------------|-------------|
| 2   | 00-05-30-00-9d-d2 to 00-05-30-00-9d-de | JAB064605a2 |
| 5   | 00-05-30-00-64-be to 00-05-30-00-64-c2 | JAB06350B1R |
| 6   | 00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd | JAB06350B1R |
| 8   | 00-05-30-01-37-7a to 00-05-30-01-37-fe | JAB072705ja |
| 9   | 00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 | JAB06280ae9 |

\* this terminal session

If the status is OK or active, you can continue with your configuration (see [Chapter 7, “Managing Modules”](#)).

## Configuring Date and Time

Switches in the Cisco MDS 9000 Family use Coordinated Universal Time (UTC), which is the same as Greenwich Mean Time (GMT). To change the default time on the switch, issue the **clock** command from EXEC mode.

```
switch# clock set <HH:MM:SS> <DD> <Month in words> <YYYY>
```

For example:

```
switch# clock set 12:07:50 23 September 2002
Mon Sep 23 12:07:50 UTC 2002
```

*HH* represents hours in military format (15 for 3 p.m.), *MM* is minutes (58), *SS* is seconds (09), *DD* is the date (02), *Month* is the month in words (August), and *YYYY* is the year (2002).



### Note

The **clock** command changes are saved across system resets.

## Configuring the Time Zone

You can specify a time zone for the switch.

To specify the local time without the daylight savings feature, follow these steps:

|        | Command                                                                                                                                                                                              | Purpose                                                                                                                                                                                                                 |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                              | Enters configuration mode.                                                                                                                                                                                              |
| Step 2 | switch(config)# <b>clock timezone</b> <timezone name> <-23 to 23 hours offset from UTC time> <0 to 50 minutes offset from UTC><br><br>Example:<br><br>switch(config)# <b>clock timezone PST -8 0</b> | Sets the time zone with a specified name, specified hours, and specified minutes.<br><br>This example sets the time zone to Pacific Standard Time (PST) and offsets the UTC time by negative eight hours and 0 minutes. |
| Step 3 | switch(config)# <b>exit</b><br>switch#                                                                                                                                                               | Returns to EXEC mode.                                                                                                                                                                                                   |
| Step 4 | switch# <b>show clock</b>                                                                                                                                                                            | Verifies the time zone configuration.                                                                                                                                                                                   |
| Step 5 | switch# <b>show run</b>                                                                                                                                                                              | Displays changes made to the time zone configuration along with other configuration information.                                                                                                                        |

## Setting the Daylight Saving Time Adjustment

Following U.S. standards, you can have the switch advance the clock one hour at 2:00 a.m. on the first Sunday in April and move back the clock one hour at 2:00 a.m. on the last Sunday in October. You can also explicitly specify the start and end dates and times and whether or not the time adjustment recurs every year.

To enable the daylight saving time clock adjustment according to the U.S. rules, follow these steps:

|        | Command                                                                                                                                                                                                                      | Purpose                                                                                                                                                                          |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                                      | Enters configuration mode.                                                                                                                                                       |
| Step 2 | switch(config)# <b>clock timezone</b> <i>timezone_name hour_offset_from_UTC minute_offset_from_UTC</i><br><br>Example:<br><br>switch(config)# <b>clock timezone PST -8 0</b><br><br>switch(config)# <b>no clock timezone</b> | Offsets the time zone as specified.<br><br>This example set the Pacific standard offset time as negative 8 hours and 0 minutes.<br><br>Disables the timezone adjustment feature. |



|        | Command                                                                                                                                                                                                                                                                                           | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 3 | <pre>switch(config)# clock summer-time daylight_timezone_name start_week start_day start_month start_time end_week end_day end_month end_time daylight_offset_inminutes</pre> <p>Example:</p> <pre>switch(config)# clock summer-time PDT 1 Sun Apr 02:00 5 Sun Oct 02:00 60 switch(config)#</pre> | <p>Sets the daylight savings time for a specified time zone.</p> <p>The start and end values are as follows:</p> <ul style="list-style-type: none"> <li>• week ranging from 1 through 5</li> <li>• day ranging from Sunday through Saturday</li> <li>• month ranging from January through December</li> </ul> <p>The daylight offset ranges from 1 through 1440 minutes which are added to the start time and deleted time from the end time.</p> <p>This example adjusts the daylight savings time for the Pacific daylight time by 60 minutes starting the first Sunday in April at 2 a.m. and ending the last Sunday in October at 2 a.m.</p> |
|        | <pre>switch(config)# no clock summer-time</pre>                                                                                                                                                                                                                                                   | Disables the daylight saving time adjustment feature.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Step 4 | <pre>switch(config)# exit switch#</pre>                                                                                                                                                                                                                                                           | Returns to EXEC mode.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Step 5 | <pre>switch# show clock</pre>                                                                                                                                                                                                                                                                     | Verifies the time zone configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## NTP Configuration

A Network Time Protocol (NTP) server provides a precise time source (radio clock or atomic clock) to synchronize the system clocks of network devices. NTP is transported over User Datagram Protocol UDP/IP. All NTP communications use UTC. An NTP server receives its time from a reference time source, such as a radio clock or atomic clock, attached to the time. NTP distributes this time across the network.

In a large enterprise network, having one time standard for all network devices is critical for management reporting and event logging functions when trying to correlate interacting events logged across multiple devices. Many enterprise customers with extremely mission-critical networks maintain their own stratum-1 NTP source.

Time synchronization happens when several frames are exchanged between clients and servers. The switches in client mode know the address of one or more NTP servers. The servers act as the time source and receive client synchronization requests.

By configuring an IP address as a peer, the switch will obtain and provide time as required. The peer is capable of providing time on its own and is capable of having a server configured. If both these instances point to different time servers, your NTP service will be more reliable. Thus, even if the active server link is lost, you can still maintain the right time due to the presence of the peer.



### Tip

If an active server fails, a configured peer helps in providing the NTP time. Provide a direct NTP server association and configure a peer to ensure backup support if the active server fails.

If you only configure a peer, the most accurate peer takes on the role of the NTP server and the other peer(s) act as a peer(s). Both machines end at the right time if they have the right time source or if they point to the right NTP source.

To configure NTP in a server association, follow these steps:

|        | Command                                                                                                                                                                                       | Purpose                                                                                                                                                                                                    |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                       | Enters configuration mode.                                                                                                                                                                                 |
| Step 2 | switch(config)# <b>ntp server 10.10.10.10</b><br>switch(config)#                                                                                                                              | Forms a server association with a server.                                                                                                                                                                  |
| Step 3 | switch(config)# <b>ntp peer 10.20.10.0</b><br>switch(config)#                                                                                                                                 | Forms a peer association with a peer. You can specify multiple associations.                                                                                                                               |
| Step 4 | switch(config)# <b>exit</b><br>switch#                                                                                                                                                        | Returns to EXEC mode.                                                                                                                                                                                      |
| Step 5 | switch# <b>copy running-config startup-config</b>                                                                                                                                             | Saves your configuration changes to NVRAM.<br><br><b>Tip</b> This is one instance where you can save the configuration as a result of an NTP configuration change. You can issue this command at any time. |
| Step 6 | switch# <b>show ntp peers</b><br><br>-----<br>Peer IP Address                    Serv/Peer<br>-----<br>10.20.10.2                          Server<br>10.20.10.0                          Peer | Displays the configured server and peer associations.<br><br><b>Note</b> A domain name will be resolved only when you have a DNS server configured.                                                        |

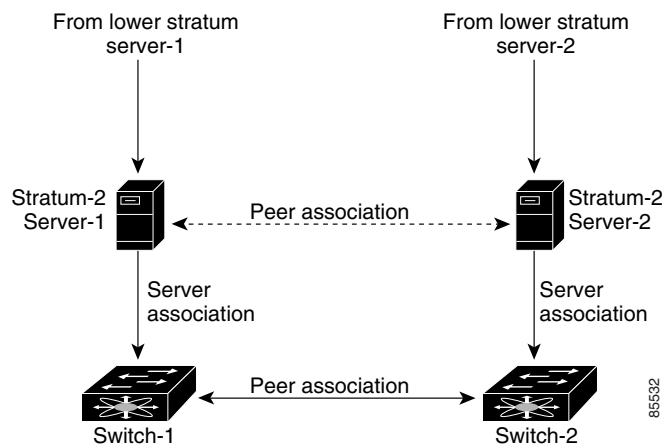
## NTP Configuration Guidelines

The following guidelines apply to all NTP configurations:

- You should have a peer association with another switch only when you are sure that your clock is reliable (which means that you are a client of a reliable NTP server).
- Though a peer configured alone, will be the most accurate peer taking on the role of a server, the configured peer should be used more as a back-up support. If more than one server is present, you can have several switches point to one server, and the remaining to the another server, and then configure peer association between these two sets. This forces the clock more reliable.
- If you only have one server, it's better for all the switches have a client association with that server.

If the network is configured robustly, even a server down time will not affect well-configured switches in the network. [Figure 4-3](#) displays a network with two NTP stratum 2 servers and two switches.

**Figure 4-3 NTP Peer and Server Association**



In this configuration, the switches were configured as explained below:

- Stratum 2 Server 1
  - IP address -10.10.10.10
  - Stratum-2 Server-2
  - IP address -10.10.10.9
- Switch 1
  - Switch ip address -10.10.10.1
- NTP Configuration
  - NTP server 10.10.10.10
  - NTP peer 10.10.10.2
- Switch 2
  - Switch ip address -10.10.10.2
  - NTP Configuration
  - NTP server 10.10.10.9
  - NTP peer 10.10.10.1

# Configuring the Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management port, but first you must configure some IP parameters (IP address, subnet mask) so that the switch is reachable. You can manually configure the management port interface from the CLI.



## Note

Before you begin to configure the management port interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

To obtain remote management access using Telnet (CLI) or SNMP (GUI), follow these steps:

|        | Command                                                                                                        | Command                                                                                                                                                                       |
|--------|----------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                                                              | Enters configuration mode. You can also abbreviate the command to <b>config t</b> .<br><br>The switch(config)# prompt indicates that you are in configuration mode.           |
| Step 2 | switch(config)# <b>interface type interface_string</b><br>Examples:<br>switch(config)# <b>interface mgmt 0</b> | Enters the interface configuration mode on the specified interface.<br><br>You can use the management Ethernet interface on the switch to configure the management interface. |
| Step 3 | switch(config)# <b>ip address 1.1.1.0 255.255.255.0</b>                                                        | Enters the IP address and IP subnet mask for the interface specified in Step 2.                                                                                               |
| Step 4 | switch(config-if)# <b>no shutdown</b>                                                                          | Enables the interface.                                                                                                                                                        |
| Step 5 | switch(config-if)# <b>exit</b>                                                                                 | Returns to configuration mode.                                                                                                                                                |
| Step 6 | switch(config)# <b>ip default-gateway 1.1.1.1</b>                                                              | Configures the default gateway address.                                                                                                                                       |

In some cases, a switch interface might be administratively shut down. You can check the status of an interface at any time by using the **show interface mgmt 0** command.

The management port (mgmt 0) is autosensing and operates as full duplex mode and 100 Mbps speed. The speed and mode cannot be configured.

When you try to shutdown a management interface (mgmt 0), a follow-up message confirms your action before performing the operation. You can use the **force** option to bypass this confirmation. The following example shuts down the interface without using the **force** option:

```
switch# config t
switch(config-if)# shutdown
Shutting down this interface will drop all telnet sessions.
Do you wish to continue (y/n)? y
```

The following example shuts down the interface using the **force** option:

```
switch# config t
switch(config-if)# shutdown force
```



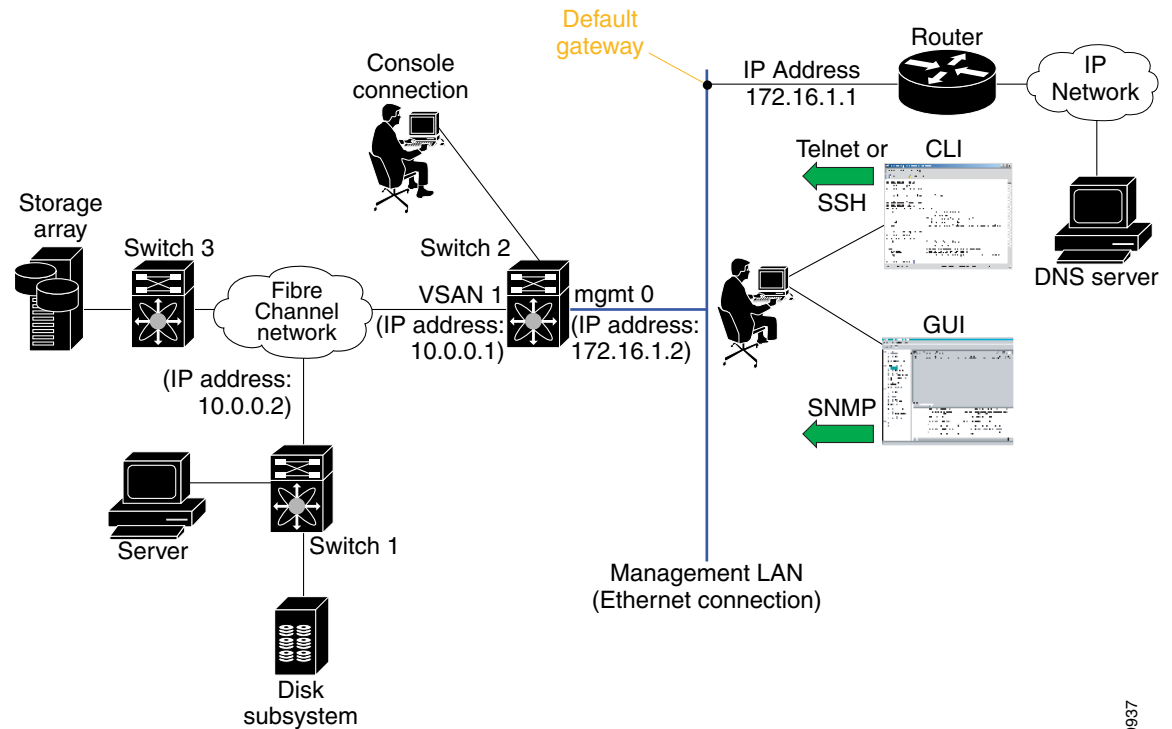
## Note

You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

## Configuring Default Gateways

The supervisor module sends IP packets with unresolved destination IP addresses to the default gateway (see Figure 4-4).

**Figure 4-4** Default Gateway



79937

To configure the IP address of the default gateway, follow these steps:

|        | Command                                                                       | Purpose                               |
|--------|-------------------------------------------------------------------------------|---------------------------------------|
| Step 1 | switch# <code>config t</code>                                                 | Enters configuration mode.            |
| Step 2 | switch(config)# <code>ip default-gateway 172.16.1.1</code><br>switch(config)# | Configures the 172.16.1.1 IP address. |

# Disabling a Telnet Server

The Telnet server is enabled by default on all switches in the Cisco MDS 9000 Family. If you require a secure SSH connection, you need to disable the default Telnet connection and then enable the SSH connection (see the [“Enabling SSH Service” section on page 16-27](#)).

**Note**

For information on connecting a terminal to the supervisor module console port, refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.

Make sure the terminal is connected to the switch and that the switch and terminal are both powered on. To allow Telnet connections to the switch, follow these steps:

|        | Command                                                   | Purpose                                                                                           |
|--------|-----------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                   | Enters configuration mode.                                                                        |
| Step 2 | switch(config)# <b>no telnet server enable</b><br>updated | Disables the Telnet server.                                                                       |
|        | switch(config)# <b>telnet server enable</b><br>updated    | Enables the Telnet server if you wish to return a Telnet connection from a secure SSH connection. |

**Tip**

A maximum of 16 sessions are allowed in any Cisco MDS 9216 Switch or in any switch in the Cisco MDS 9500 Series.

# Working with Configuration Files

This section describes how to work with configuration files and has the following topics:

- [Guidelines for Creating and Using Configuration Files, page 4-23](#)
- [Viewing Configuration Files, page 4-23](#)
- [Downloading Configuration Files to the Switch, page 4-23](#)
- [Saving the Configuration, page 4-26](#)
- [Copying Files, page 4-27](#)

## Guidelines for Creating and Using Configuration Files

Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration so that they have identical module and port configurations.

## Viewing Configuration Files

To view the running configuration file, use the **show running-config** command:

```
switch# show running-config
Building Configuration ...
 interface port-channel 98
interface fc1/1
 interface fc1/2
interface mgmt0
ip address 172.22.95.112 255.255.255.0
no shutdown
vsan database
vsan 2
clock summer-time Pacific 1 Sun Apr 02:00 5 Sun Oct 02:00 60
switchname switch112
```

To view the startup configuration file, use the **show startup-config** command:

```
switch# show startup-config
 interface port-channel 98
 interface fc1/1
channel-group 98 force
no shutdown
 interface mgmt0
ip address 172.22.95.112 255.255.255.0
boot system system-237; ep-41
boot kickstart boot-237 ep-41
ip domain-name cisco.com
```

## Downloading Configuration Files to the Switch

You can configure a switch in the Cisco MDS 9000 Family by using configuration files you create or download from another switch. In addition, you can store configuration files on a bootflash device on the supervisor module and you can configure the switch using a configuration stored on an external CompactFlash disk.

Before you begin downloading a configuration file using a remote server, do the following:

- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router or default gateway to route traffic between subnets. Check connectivity to the remote server using the **ping** command.
- Ensure the configuration file to be downloaded is in the correct directory on the remote server.
- Ensure that the permissions on the file are set correctly. Permissions on the file should be set to world-read.

**Note**

See the [“Copying Files” section on page 4-27](#).

## From a Remote Server

To configure a switch in the Cisco MDS 9000 Family using a configuration file downloaded from a remote server using TFTP, FTP, SCP, or SFTP, follow these steps:

- 
- Step 1** Log into the switch through the console port or through a Telnet session.
- Step 2** Configure the switch using the configuration file downloaded from the remote server using the **copy <scheme> :// <server address> system:running-config** command.

The *scheme* is TFTP, FTP, SCP, or SFTP.

The configuration file downloads and the commands are executed as the file is parsed line by line.

---

Use the following command to download a configuration file from a remote server to the running configuration.

```
switch# copy <scheme>://<url> system:running-config
```

Use the following command to download a configuration file from a remote server to the startup configuration.

```
switch# copy <scheme>://<url> nvram:startup-config
```



## From an External Flash (slot0:)



### Note

The physical media must be inserted into slot0: after you log into the switch.

To configure a switch in the Cisco MDS 9000 Family using a configuration file stored on an external CompactFlash disk, follow these steps:

- 
- Step 1** Log into the switch through the console port or through a Telnet session.
  - Step 2** Locate the configuration file using the **cd** and **dir** commands. (See the “Copying Files” section on [page 4-27](#).)
  - Step 3** Configure the switch using the configuration file stored on the external CompactFlash disk using the **copy <source file> system:running-config** command.  
The commands are executed as the file is parsed line by line.
- 

Use the following command to download a configuration file from an external CompactFlash to the running configuration:

```
switch copy slot0:dns-config.cfg system:running-config
```

Use the following command to download a configuration file from an external CompactFlash to the startup configuration:

```
switch copy slot0:dns-config.cfg nvram:startup-config
```

## To a Remote Server

To save a configuration file to a remote server like TFTP, FTP, SCP, or SFTP, follow these steps:

- 
- Step 1** Log into the switch through the console port or through a Telnet session.
  - Step 2** Save the configuration using the **copy system: running-config <scheme> :// <url>** command.  
Scheme can be TFTP, FTP, SCP, or SFTP.
  - Step 3** Specify the IP address or host name of the remote server and the name of the file to download.  
The configuration file is saved to the remote server.
- 

Use the following command to save a running configuration file to a remote server:

```
switch# copy system:running-config <scheme>://<url>
```

Use the following command to save a startup configuration file to a remote server:

```
switch# copy nvram:startup-config <scheme>://<url>
```

## To an External CompactFlash Disk

To save a configuration file on an external CompactFlash disk, follow these steps:

- 
- Step 1** Log into the switch through the console port or through a Telnet session.
  - Step 2** Locate the configuration file using the **cd** and **dir** commands. (See the “Copying Files” section on [page 4-27](#).)
  - Step 3** Save the configuration file using the **copy system:running-config <source file>** command.  
The configuration file is saved to the CompactFlash disk.
- 

Use the following command to save a running configuration file to an external CompactFlash disk:

```
switch# copy system:running-config slot0:dns-config.cfg
```

Use the following command to save a startup configuration file to an external CompactFlash disk:

```
switch# copy system:startup-config slot0:dns-config.cfg
```

## Saving the Configuration

After you have created a configuration, you save the configuration using the following **copy** command:

```
switch# copy system:running-config nvram:startup-config
```

The **copy running-config startup-config** command is an alias to the previous command and is used frequently throughout this guide.

# Copying Files

The syntax for the **copy** command follows and is explained in [Table 4-1](#).

```
switch# copy <scheme>://<username@><server>/<file name>
<scheme>://<username@><server>/<file name>
```

**Table 4-1** *copy Command Syntax*

| Scheme            | Server                                                                                         | File Name                         |
|-------------------|------------------------------------------------------------------------------------------------|-----------------------------------|
| bootflash         | sup-active<br>sup-standby<br>sup-1 or module-5<br>sup-2 or module-6<br>sup-local<br>sup-remote | User-specified                    |
| slot0             | —                                                                                              | User-specified                    |
| volatile          | —                                                                                              | User-specified                    |
| nvr               | —                                                                                              | startup-config or snapshot-config |
| system            | —                                                                                              | running-config                    |
| tftp <sup>1</sup> | IP address or DNS name                                                                         | User-specified                    |
| ftp               |                                                                                                |                                   |
| scp (secure copy) |                                                                                                |                                   |
| sftp              |                                                                                                |                                   |
| core              | slot-number                                                                                    | Process identifier number         |

1. When downloading and uploading files, a TFTP limitation restricts a TFTP client to a 32 MB file size and some TFTP servers to a 16 MB file size.

- This example shows how to copy a file from the active supervisor module's (sup-1 in slot 5) bootflash to the standby supervisor module's (sup-2 in slot 6) bootflash.

```
switch# copy bootflash:system_image bootflash://sup-2/system_image
```

- This example shows how to overwrite the contents of an existing configuration in NVRAM.

```
switch# copy nvram:snapshot-config nvram:startup-config
Warning: this command is going to overwrite your current startup-config.
Do you wish to continue? {y/n} [y] y
```

- This example shows how to create a running configuration copy in bootflash.

```
switch# copy system:running-config bootflash:my-config
```

- This example shows how to copy a system image file from the TFTP server to bootflash.

```
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
```

- This example shows how to copy a script file from the SFTP server to volatile.

```
switch# copy sftp://172.16.10.100/myscript.txt volatile:myscript.txt
```



## Note

Use the **show version image** command to verify if the downloaded images are valid.

## Backing up the Current Configuration

Before installing or migrating to any software configuration, back up the startup configuration.

- This example shows how to create a snapshot of the startup configuration in a predefined location on the switch (binary file).

```
switch# copy nvram:startup-config nvram:snapshot-config
```

- This example shows how to backup the startup configuration copy in the bootflash: file system (ASCII file).

```
switch# copy nvram:startup-config bootflash:my-config
```

- This example shows how to backup the startup configuration to the TFTP server (ASCII file).

```
switch# copy nvram:startup-config copy tftp://172.16.10.100/my-config
```

- This example shows how to backup the running configuration to the bootflash: file system (ASCII file).

```
switch# copy system:running-config bootflash:my-config
```

## Rolling Back to a Previous Configuration

All switch configurations reside in the internal bootflash: file system. If your internal bootflash: file system is corrupted, you could potentially lose your configuration. Save and back up your configuration file periodically.

- This example shows how to roll back to a snapshot copy of a previously saved running configuration (binary file).

```
switch# copy nvram:snapshot-config nvram:startup-config
```



### Note

You can issue a rollback command only when a snapshot is already created. Otherwise, you will receive the `No snapshot-config found` error message.

- This example shows how to roll back to a configuration copy that was previously saved in the bootflash: file system (ASCII file).

```
switch# copy bootflash:my-config nvram:startup-config
```

- This example shows how to roll back to a configuration copy that was previously saved in a TFTP server (ASCII file).

```
switch# copy tftp://172.16.10.100/my-config nvram:startup-config
```



### Note

Each time a **copy running-config startup-config** command is issued a binary file is created and the ASCII file is updated. A valid binary configuration file reduces the overall boot time significantly. A binary file cannot be uploaded, but its contents can be used to overwrite the existing startup configuration. The **write erase** command clears the binary file to match the new system image.

## Restoring the Configured Redundancy Mode



### Tip

If you have configured the **combined** mode as the redundancy mode for power supplies on a Cisco MDS 9509 switch, exert care when using the sequence of the **write erase** and **reload** commands before rolling back to a saved configuration.

As a result of issuing the **write erase** command and the **reload** command, you restore the switch settings to their factory defaults. This sequence also restores the redundancy mode setting for the power supplies back to the **redundant** mode (default).

Depending on the types of power supplies, the input voltage, and the number of modules (line cards) in the chassis, the redundancy mode may prevent the line cards from being powered on after a system reboot (see the [“Configuring Power Supplies”](#) section on page 8-6).

If you use this sequence, the commands that apply to the powered down line cards will not be enforced on the switch (and will not be part of its running configuration).

When using the sequence of the **write erase** and **reload** commands before rolling back to a saved configuration, follow these steps:

- 
- Step 1** Manually change the **redundant** mode configuration to **combined** mode, if originally configured as such.
  - Step 2** Wait until all modules are back online—the module status displays `ok` in response to the **show module** command.
  - Step 3** Rollback to the saved configuration (see the [“Rolling Back to a Previous Configuration”](#) section on page 4-28).
- 

## Accessing Remote File Systems

To access contents of the standby supervisor module (remote), follow these steps:

- 
- Step 1** Verify if the standby supervisor module has sufficient space for new image files.

```
switch# dir bootflash://sup-remote
 12198912 Aug 27 17:21:10 2003 bootflash:boot-39a
 12198912 Aug 27 16:29:18 2003 bootflash:m9500-sflek9-kickstart-mzg.1.3.0.39a.bin
 1921922 Sep 14 19:58:12 2003 aOldImage
 1864931 Apr 29 12:41:50 2003 bOldImage
 1864931 Apr 29 12:41:59 2003 dplug2
 12288 Apr 18 20:23:11 2003 lost+found/
 12097024 Nov 21 16:34:18 2003 m9500-sflek9-kickstart-mz.1.3.1.1.bin
 41574014 Nov 21 16:34:47 2003 m9500-sflek9-mz.1.3.1.1.bin
 1024 Oct 28 20:24:59 2003 newer-fs/
 2021518 Oct 11 15:49:41 2003 plugin-69a

Usage for bootflash://sup-remote
 102081536 bytes used
 82478080 bytes free
 184559616 bytes total
```

- Step 2** Delete files, if required, to make more space for the new image files.

```
switch# del aOldImage
switch# del aOldImage
```

## Downgrading from a Higher Release

When downgrading any switch in the Cisco MDS 9000 Family, avoid using the **reload** command:

**Tip**

Use the **install all** command to gracefully reload the switch and handle configuration conversions.

For example, to revert to Release 1.0(4) or 1.0(3a) from Release 1.x, follow these steps:

- Step 1** Save the configuration using the **copy running-config startup-config** command to save the new configuration into nonvolatile storage.
- Step 2** Issue the **install all** command to reload the switch.
- Refer to the *Cisco MDS 9000 Family Configuration Guide* for further information on the **install all** command.

## Deleting Files

To delete files on a Flash device, follow these steps:

|        | Command                                  | Purpose                         |
|--------|------------------------------------------|---------------------------------|
| Step 1 | switch# <b>delete</b> [device:] filename | Deletes files from a directory  |
| Step 2 | switch# <b>dir</b> [device:] [filename]  | Verifies the files are deleted. |

- This example shows how to delete a file from the bootflash: directory (assuming you are already in the bootflash: directory):

```
switch# delete dns_config.cfg
```

- This example shows how to delete a file from an external CompactFlash (slot0):

```
switch# delete slot0:dns_config.cfg
```

- This example deletes the file named test from the Flash card inserted in slot 0.

```
switch# delete slot0:test
Delete slot0:test? [confirm]
```

- This example deletes the entire my-dir directory and all its contents:

```
switch# delete bootflash:my-dir
```

# Configuring Console Settings

A console port is an asynchronous serial port that enables switches in the Cisco MDS 9000 Family to be set up for initial configuration through a standard RS-232 port with an RJ-45 connector. Any device connected to this port must be capable of asynchronous transmission. Connection to a terminal requires a terminal emulator to be configured as 9600 baud, 8 data bits, 1 stop bit, no parity.



## Caution

The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

To configure the console port parameters from the console terminal, follow these steps:

|        | Command                                                        | Command                                                                                                                                                                                                                                                                      |
|--------|----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#              | Enters configuration mode.                                                                                                                                                                                                                                                   |
| Step 2 | switch(config)# <b>line console</b><br>switch(config-console)# | Enters the line console configuration mode.                                                                                                                                                                                                                                  |
| Step 3 | switch(config-console)# <b>speed</b><br>9600                   | Configures the port speed for the serial console. The default console baud rate is 9600 baud. The valid range is between 110 and 115,200 bps (110, 150, 300, 600, 1200, 2400, 4800, 9600, 19200, 28800, 38400, 57600, 115200). Be sure to specify one of these exact values. |
| Step 4 | switch(config-console)# <b>databits</b> 8                      | Configures the data bits for the console connection. The default is 8 data bits and the valid range is between 5 and 8 data bits.                                                                                                                                            |
| Step 5 | switch(config-console)# <b>stopbits</b> 1                      | Configures the stop bits for the console connection. The default is 1 stop bit and the valid values are 1 or 2 stop bits.                                                                                                                                                    |
| Step 6 | switch(config-console)# <b>parity</b><br>none                  | Configures the parity for the console connection. The default is no parity and the valid values are even or odd parity.                                                                                                                                                      |

## Verifying the Console Configuration

Use the **show line console** command to verify the configured console settings. This command also displays problems that may have occurred along with the other registration statistics.

```
switch# show line console
line Console:
 Speed: 9600 bauds
 Databits: 8 bits per byte
 Stopbits: 1 bit(s)
 Parity: none
 Modem In: Enable
 Modem Init-String -
 default : ATE0Q1&D2&C1S0=1\015
 Statistics: tx:12842 rx:366 Register Bits:RTS|CTS|DTR|DSR|CD|RI
```

## Configuring COM1 Settings

A COM1 port is a RS-232 port with a DB-9 interface that enables you to connect to an external serial communication device such as a modem. Connection to a terminal requires terminal emulator to be configured as 9600 baud, 8 data bits, 1 stop bit, no parity.

To configure the COM1 port parameters, follow these steps:

|        | Command                                                  | Command                                                                                                                                                                                                                                                                                                                                                                           |
|--------|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#        | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                                        |
| Step 2 | switch(config)# <b>line com1</b><br>switch(config-com1)# | Enters the COM1 port configuration mode.                                                                                                                                                                                                                                                                                                                                          |
| Step 3 | switch(config-com1)# <b>speed 9600</b>                   | Configures the port speed for the COM1 connection. The default console baud rate is 9600 baud. The valid range is between 110 and 115,200 bps (110, 150, 300, 600, 1200, 2400, 4800, 9600, 19200, 28800, 38400, 57600, 115200). Be sure to specify one of these exact values.<br><br><b>Note</b> This configuration depends on the incoming speed of the modem connected to COM1. |
| Step 4 | switch(config-com1)# <b>databits 8</b>                   | Configures the data bits for the COM1 connection. The default is 8 data bits and the valid range is between 5 and 8 data bits.                                                                                                                                                                                                                                                    |
| Step 5 | switch(config-com1)# <b>stopbits 1</b>                   | Configures the stop bits for the COM1 connection. The default is 1 stop bits and the valid values are 1 or 2 stop bits.                                                                                                                                                                                                                                                           |
| Step 6 | switch(config-com1)# <b>parity none</b>                  | Configures the parity for the COM1 connection. The default is no parity and the valid values are even or odd parity.                                                                                                                                                                                                                                                              |
| Step 7 | switch(config-com1)# <b>no flowcontrol hardware</b>      | Disables hardware flowcontrol. By default, hardware flowcontrol is enabled on all switches in the Cisco 9000 Family. When enabled, this option is useful in protecting data loss at higher baud rates.<br><br><b>Note</b> This option is only available through the COM1 port.                                                                                                    |

## Verifying the COM1 Configuration

Use the **show line com1** command to verify the configured COM1 settings. This command also displays problems that may have occurred along with the other registration statistics.

```
switch# show line com1
line Aux:
 Speed: 9600 bauds
 Databits: 8 bits per byte
 Stopbits: 1 bit(s)
 Parity: none
 Modem In: Enable
 Modem Init-String -
 default : ATE0Q1&D2&C1S0=1\015
 Statistics: tx:17 rx:0 Register Bits:RTS|DTR
```



# Configuring Modem Connections

Modems can only be configured if you are connected to the console or COM1 ports. A modem connection to a switch in the Cisco MDS 9000 Family does not affect switch functionality.



## Note

If you plan on connecting a modem to the console port or the COM1 port of a switch in the Cisco MDS 9000 Family, refer to the *Cisco MDS 9216 Switch Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*. COM1 ports are not available on switches in the Cisco MDS 9100 Series, refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

## Guidelines to Configure Modems

The following guidelines apply to modem configurations:

- Connect the modem before attempting to configure the modem.
- Do not connect a modem to the console port while the system is booting. Follow the procedure specified in the “[Initializing a Modem in a Powered on Switch](#)” section on page 4-36.
- The following Cisco modems have been verified to work in the SAN-OS environment:
  - MultiTech MT2834BA (<http://www.multitech.com/PRODUCTS/Families/MultiModemII/>)
  - Hayes Accura V.92 (<http://www.hayesmicro.com/Products/accura-prod-v92.htm>)
- We recommend you use the COM1 port to connect the modem from a Cisco MDS 9216 switch or from any director in the Cisco MDS 9500 Family.

## Enabling Modem Connections

To configure a modem connection through the COM1 port, follow these steps:

|        | Command                                                  | Command                                                     |
|--------|----------------------------------------------------------|-------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#               | Enters configuration mode.                                  |
| Step 2 | switch(config)# <b>line com1</b><br>switch(config-com1)# | Enters the COM1 port configuration mode.                    |
| Step 3 | switch(config-com1)# <b>no modem in</b>                  | Disables the current modem from executing its functions.    |
|        | switch(config-com1)# <b>modem in</b>                     | Enables (default) the COM1 port to only connect to a modem. |

To configure a modem connection through the console port, follow these steps:

|        | Command                                                        | Command                                     |
|--------|----------------------------------------------------------------|---------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                     | Enters configuration mode.                  |
| Step 2 | switch(config)# <b>line console</b><br>switch(config-console)# | Enters the console port configuration mode. |

## Step 3

| Command                                          | Command                                                            |
|--------------------------------------------------|--------------------------------------------------------------------|
| <code>switch(config-console)# modem in</code>    | Enables the console port to only connect to a modem.               |
| <code>switch(config-console)# no modem in</code> | Disables (default) the current modem from executing its functions. |

## Configuring the Initialization String

Switches in the Cisco MDS 9500 Family and the Cisco MDS 9216 switch have a default initialization string (`ATE0Q1&D2&C1S0=1\015`) to detect connected modems. The default string detects connected modems only the modem is supported by Cisco systems. The default string contents are as follows:

- AT—Attention
- E0 (required)—No echo
- Q1—Result code on
- &D2—Normal data terminal ready (DTR) option
- &C1—Enable tracking the state of the data carrier.
- S0=1—Pick up after one ring
- \015 (required)—carriage return in octal

You may retain the default string or change it to another string (80 character limit) using the **user-input** option. This option is provided if you prefer to use a modem that is not supported or tested by Cisco systems. If you change the string, the changes you make are permanent and remain in effect unless you change them again. Rebooting the system or restarting the CLI does not change the modem initialization string. The switch is not affected even if the modem is not functioning.

**Tip**

We recommend you use the default initialization string. If the required options are not provided in the user-input string, the initialization string is not processed.

The modem initialization string usage depends on the modem state when the switch boots:

- If the modem is already attached to the switch during boot-up, the default initialization string is written to the modem (see the [“Configuring the Default Initialization String”](#) section on page 4-35).
- If the modem is not attached to the switch during boot-up, then attach the modem as outlined in the *Cisco MDS 9000 Family Hardware Installation Guide* (depending on the product), and follow the procedure provided in this section (see the [“Configuring a User-Specified Initialization String”](#) section on page 4-35).

**Note**

You can perform the configuration specified in this section only if you are connected to the console port or the COM1 port.

## Configuring the Default Initialization String

To configure the default initialization string through the COM1 port, follow these steps:

|        | Command                                                  | Command                                                |
|--------|----------------------------------------------------------|--------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#               | Enters configuration mode.                             |
| Step 2 | switch(config)# <b>line com1</b><br>switch(config-com1)# | Enters the COM1 port configuration mode.               |
| Step 3 | switch(config-com1)# <b>modem init-string default</b>    | Writes the default initialization string to the modem. |

To configure the default initialization string through the console port, follow these steps:

|        | Command                                                     | Command                                                |
|--------|-------------------------------------------------------------|--------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                  | Enters configuration mode.                             |
| Step 2 | switch(config)# <b>line com1</b><br>switch(config-console)# | Enters the console port configuration mode.            |
| Step 3 | switch(config-console)# <b>modem init-string default</b>    | Writes the default initialization string to the modem. |

## Configuring a User-Specified Initialization String

To configure a user-specified initialization string through the COM1 port, follow these steps:

|        | Command                                                                                 | Command                                                                                                                                                                     |
|--------|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                                       | Enters configuration mode.                                                                                                                                                  |
| Step 2 | switch(config)# <b>line com1</b><br>switch(config-com1)#                                | Enters the COM1 port configuration mode.                                                                                                                                    |
| Step 3 | switch(config-com1)# <b>modem set-string user-input ATE0Q1&amp;D2&amp;C1S0=3\015</b>    | Assigns the user-specified initialization string to its corresponding profile.<br><br><b>Note</b> You must first set the user-input string, before initializing the string. |
|        | switch(config-com1)# <b>no modem set-string user-input ATE0Q1&amp;D2&amp;C1S0=3\015</b> | Deletes the configured initialization string.                                                                                                                               |
| Step 4 | switch(config-com1)# <b>modem init-string user-input</b>                                | Writes the user-specified initialization string to the modem.                                                                                                               |

To configure a user-specified initialization string through the console port, follow these steps:

|        | Command                                                     | Command                                     |
|--------|-------------------------------------------------------------|---------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#           | Enters configuration mode.                  |
| Step 2 | switch(config)# <b>line com1</b><br>switch(config-console)# | Enters the console port configuration mode. |

|        | Command                                                                                    | Command                                                                                                                                                                     |
|--------|--------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 3 | switch(config-console)# <b>modem set-string user-input ATE0Q1&amp;D2&amp;C1S0=3\015</b>    | Assigns the user-specified initialization string to its corresponding profile.<br><br><b>Note</b> You must first set the user-input string, before initializing the string. |
|        | switch(config-console)# <b>no modem set-string user-input ATE0Q1&amp;D2&amp;C1S0=3\015</b> | Deletes the configured initialization string.                                                                                                                               |
| Step 4 | switch(config-console)# <b>modem init-string user-input</b>                                | Writes the user-specified initialization string to the modem.                                                                                                               |

## Initializing a Modem in a Powered on Switch

When a switch is already powered on and the modem is later connected to either the console port or the COM1 port, you can initialize the modem using the **modem connect line** command in EXEC mode. You can specify the **com1** option if the modem is connected to the COM1 port, or the **console** option if the modem is connected to the console.

To connect a modem to a switch that is already powered on, follow these steps.

- 
- |        |                                                                                                                                                                   |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Wait till the system has completed the boot sequence and the system image is running.                                                                             |
| Step 2 | Connect the modem to the switch as specified in the <i>Cisco MDS 9216 Switch Hardware Guide</i> or the <i>Cisco MDS 9500 Series Hardware Installation Guide</i> . |
| Step 3 | Initialize the modem using the <b>modem connect line</b> command in EXEC mode.                                                                                    |
-

# Configuring CDP

The Cisco Discovery Protocol (CDP) is an advertisement protocol used by Cisco devices to advertise itself to other Cisco devices in the same network. CDP runs on the data link layer and is independent of layer 3 protocols. Cisco devices that receive the CDP packets cache the information to make it is accessible over CLI and SNMP.

CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module. The CDP daemon is restartable and switchable. The running and startup configurations are available across restarts and switchovers.

CDP version 1 (v1) and version 2 (v2) are supported in Cisco MDS 9000 Family switches. CDP packets with any other version number are silently discarded when received.

When the interface link is established, CDP is enabled by default and three CDP packets are sent at one-second intervals. Following this, the CDP frames are sent at the globally-configured refresh interval.

To globally disable the CDP protocol, follow these steps:

|        | Command                                                                                                                  | Command                                                                                                                                                                            |
|--------|--------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                                                                        | Enters configuration mode.                                                                                                                                                         |
| Step 2 | switch(config)# <b>no cdp enable</b><br>Operation in progress. Please check global parameters<br>switch(config-console)# | Disables the CDP protocol on the switch. When CDP is disabled on an interface, one packet is sent to clear out the switch state with each of the receiving devices.                |
|        | switch(config)# <b>cdp enable</b><br>Operation in progress. Please check global parameters<br>switch(config)#            | Enables (default) the CDP protocol on the switch. When CDP is enabled on an interface, one packet is sent immediately. Subsequent packets are sent at the configured refresh time. |

To disable the CDP protocol on a specific interface, follow these steps:

|        | Command                                                                                                                        | Command                                                                                                                                                                                        |
|--------|--------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                                                                              | Enters configuration mode.                                                                                                                                                                     |
| Step 2 | switch(config)# <b>interface gigabitethernet 8/8</b><br>switch(config-if)#                                                     | Configures the Gigabit Ethernet interface for the module in slot 8 port 8.                                                                                                                     |
| Step 3 | switch(config-if)# <b>no cdp enable</b><br>Operation in progress. Please check interface parameters<br>switch(config-console)# | Disables the CDP protocol on the selected interface. When CDP is disabled on an interface, one packet is sent to clear out the switch state with each of the receiving devices.                |
|        | switch(config-if)# <b>cdp enable</b><br>Operation in progress. Please check interface parameters<br>switch(config)#            | Enables (default) the CDP protocol on the selected interface. When CDP is enabled on an interface, one packet is sent immediately. Subsequent packets are sent at the configured refresh time. |

To globally configure the refresh time interval for the CDP protocol, follow these steps:

|        | Command                                                    | Command                                                                                                            |
|--------|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#          | Enters configuration mode.                                                                                         |
| Step 2 | switch(config)# <b>cdp timer 100</b><br>switch(config)#    | Sets the refresh time interval in seconds. The default is 60 seconds and the valid range is from 5 to 255 seconds. |
|        | switch(config)# <b>no cdp timer 100</b><br>switch(config)# | Reverts the refresh time interval to the factory default of 60 seconds.                                            |

To globally configure the hold time advertised in CDP packets, follow these steps:

|        | Command                                                       | Command                                                                                                                            |
|--------|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#             | Enters configuration mode.                                                                                                         |
| Step 2 | switch(config)# <b>cdp holdtime 200</b><br>switch(config)#    | Sets the hold time advertised in CDP packets in seconds. The default is 180 seconds and the valid range is from 10 to 255 seconds. |
|        | switch(config)# <b>no cdp holdtime 200</b><br>switch(config)# | Reverts the hold time to the factory default of 180 seconds.                                                                       |

To globally configure the CDP version, follow these steps:

|        | Command                                                    | Command                                                                                          |
|--------|------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#          | Enters configuration mode.                                                                       |
| Step 2 | switch(config)# <b>cdp advertise v1</b><br>switch(config)# | Sets the CDP version to be used. The default is version 2 (v2). The valid options are v1 and v2. |
|        | switch(config)# <b>no advertise v1</b><br>switch(config)#  | Reverts the version to the factory default of v2.                                                |

## Clearing CDP Configurations

To clear CDP traffic counters for all interfaces use the **clear cdp counters** command. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces):

```
switch# clear cdp counters
switch#
```

To clear neighboring CDP entries for all interfaces, use the **clear cdp table** command. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces):

```
switch# clear cdp table interface gigabitethernet 4/1
switch#
```

## Displaying CDP Protocol Settings

Use the **show cdp** command to display CDP entries. See Examples 4-1 to 4-11.

### **Example 4-1** Displays All CDP Capable Interfaces and Parameters

```
switch# show cdp all
GigabitEthernet4/1 is up
 CDP enabled on interface
 Sending CDP packets every 60 seconds
 Holdtime is 180 seconds
GigabitEthernet4/8 is down
 CDP enabled on interface
 Sending CDP packets every 60 seconds
 Holdtime is 180 seconds
mgmt0 is up
 CDP enabled on interface
 Sending CDP packets every 100 seconds
 Holdtime is 200 seconds
```

### **Example 4-2** Displays All CDP Neighbor Entries

```
switch# show cdp entry all

Device ID:069038747(Kiowa3)
Entry address(es):
 IP Address: 172.22.92.5
Platform: WS-C5500, Capabilities: Trans-Bridge Switch
Interface: mgmt0, Port ID (outgoing port): 5/22
Holdtime: 136 sec

Version:
WS-C5500 Software, Version McpSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems

Advertisement Version: 1
```

### **Example 4-3** Displays the Specified CDP Neighbor

```
switch# show cdp entry name 0

Device ID:0
Entry address(es):
 IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 144 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
```

**Example 4-4 Displays Global CDP Parameters**

```
switch# show cdp global
Global CDP information:
 CDP enabled globally
 Sending CDP packets every 60 seconds
 Sending a holdtime value of 180 seconds
 Sending CDPv2 advertisements is enabled
```

**Example 4-5 Displays CDP Parameters for the Management Interface**

```
switch# show cdp interface mgmt 0
mgmt0 is up
 CDP enabled on interface
 Sending CDP packets every 60 seconds
 Holdtime is 180 seconds
```

**Example 4-6 Displays CDP Parameters for the Gigabit Ethernet Interface**

```
switch# show cdp interface gigabitethernet 4/1
GigabitEthernet4/1 is up
 CDP enabled on interface
 Sending CDP packets every 80 seconds
 Holdtime is 200 seconds
```

**Example 4-7 Displays CDP Neighbors (brief)**

```
switch# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
 S - Switch, H - Host, I - IGMP, r - Repeater
```

| Device ID        | Local Intrfce | Hldtme | Capability | Platform      | Port ID |
|------------------|---------------|--------|------------|---------------|---------|
| 0                | Gig4/1        | 135    | H          | DS-X9530-SF1- | Gig4/1  |
| 069038732(Kiowa2 | mgmt0         | 132    | T S        | WS-C5500      | 8/11    |
| 069038747(Kiowa3 | mgmt0         | 156    | T S        | WS-C5500      | 6/20    |
| 069038747(Kiowa3 | mgmt0         | 158    | T S        | WS-C5500      | 5/22    |

**Example 4-8 Displays CDP Neighbors (detail)**

```
switch# show CDP neighbor detail

Device ID:0
Entry address(es):
 IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 162 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full

Device ID:069038732(Kiowa2)
Entry address(es):
 IP Address: 172.22.91.5
Platform: WS-C5500, Capabilities: Trans-Bridge Switch
Interface: mgmt0, Port ID (outgoing port): 8/11
Holdtime: 132 sec
```



```

Version:
WS-C5500 Software, Version MpsSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems
Advertisement Version: 1

```

**Example 4-9 Displays the Specified CDP Neighbor (detail)**

```

switch# show cdp neighbors interface gigabitethernet 4/1 detail

Device ID:0
Entry address(es):
 IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 144 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full

```

**Example 4-10 Displays CDP Traffic Statistics for the Management Interface**

```

switch# show cdp traffic interface mgmt 0

Traffic statistics for mgmt0
Input Statistics:
 Total Packets: 1148
 Valid CDP Packets: 1148
 CDP v1 Packets: 1148
 CDP v2 Packets: 0
 Invalid CDP Packets: 0
 Unsupported Version: 0
 Checksum Errors: 0
 Malformed Packets: 0
Output Statistics:
 Total Packets: 2329
 CDP v1 Packets: 1164
 CDP v2 Packets: 1165
 Send Errors: 0

```

**Example 4-11 Displays CDP Traffic Statistics for the Gigabit Ethernet Interface**

```

switch# show cdp traffic interface gigabitethernet 4/1

Traffic statistics for GigabitEthernet4/1
Input Statistics:
 Total Packets: 674
 Valid CDP Packets: 674
 CDP v1 Packets: 0
 CDP v2 Packets: 674
 Invalid CDP Packets: 0
 Unsupported Version: 0
 Checksum Errors: 0
 Malformed Packets: 0
Output Statistics:
 Total Packets: 674
 CDP v1 Packets: 0
 CDP v2 Packets: 674
 Send Errors: 0

```





## Configuring High Availability

---

This chapter provides details on the high availability feature that is available on switches with two supervisor modules. It includes the following sections:

- [About High Availability, page 5-2](#)
- [Switchover Mechanisms, page 5-3](#)
- [Switchover Guidelines, page 5-3](#)
- [Process Restartability, page 5-4](#)
- [Synchronizing Supervisor Modules, page 5-4](#)
- [Automatically Copying Images to the Standby Supervisor, page 5-4](#)
- [Displaying HA Information, page 5-5](#)

# About High Availability

The Cisco MDS 9500 Series of multilayer directors support application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework. The high availability (HA) software framework provides for the following:

- Ensures nondisruptive software upgrade capability. See [Chapter 6, “Software Images.”](#)
- Provides redundancy for supervisor module failure by using dual supervisor modules.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Protects against link failure using the PortChannel (port aggregation) feature. See [Chapter 12, “Configuring PortChannels.”](#) This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Provides management redundancy using Virtual Router Redundancy Protocol (VRRP). See the [“Configuring VRRP” section on page 20-18](#). This feature is also available in Cisco MDS 9216 switches and in the Cisco MDS 9100 Series.
- Switchability—When the active supervisor fails, the standby supervisor, if present, takes over without disrupting storage or host traffic.

Directors in the Cisco MDS 9500 Series have two supervisor modules in the two center slots (sup-1 and sup-2). When the switch powers up and both supervisor modules are present, the supervisor module that comes up first enters the active mode and the supervisor module that comes up second enters the standby mode. If both supervisor modules come up at the same time, sup-1 becomes active. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

# Switchover Mechanisms

When the active supervisor module fails, the standby module automatically takes over. You can also issue a **system switchover** command to manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a **system switchover** is issued (switchover process has started) another switchover process cannot be started on the same switch until a stable standby supervisor module is available.



## Caution

If the supervisor modules are not in a stable state (online or powered down), a switchover will not be performed.

## HA Switchover

When a **show system redundancy status** or a **show module** command displays the HA-standby state for the standby supervisor module, an HA switchover is possible. An HA switchover has the following characteristics:

- Is stateful (nondisruptive) since control traffic is not impacted
- Does not impact data traffic since the switching modules are not impacted
- Switching modules are not reset

## Switchover Guidelines

Be aware of the following guidelines when performing a switchover:

- The **system switchover** command returns the following message when the standby supervisor is not present in the switch:

```
switch# system switchover
Failed to switchover (supervisor has no standby)
```

- You can only perform a switchover when the switch has two supervisor modules functioning in the switch. Use the **show system redundancy status** command to ensure that the system is ready to accept a switchover.
- Verify that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
```

| Mod | Ports | Module-Type                | Model           | Status     |
|-----|-------|----------------------------|-----------------|------------|
| 2   | 8     | IP Storage Services Module | DS-X9308-SMIP   | ok         |
| 5   | 0     | Supervisor/Fabric-1        | DS-X9530-SF1-K9 | active *   |
| 6   | 0     | Supervisor/Fabric-1        | DS-X9530-SF1-K9 | ha-standby |
| 8   | 0     | Caching Services Module    | DS-X9560-SMAP   | ok         |
| 9   | 32    | 1/2 Gbps FC Module         | DS-X9032        | ok         |

| Mod | Sw          | Hw    | World-Wide-Name(s) (WWN)                           |
|-----|-------------|-------|----------------------------------------------------|
| 2   | 1.3(0.106a) | 0.206 | 20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00 |
| 5   | 1.3(0.106a) | 0.602 | --                                                 |
| 6   | 1.3(0.106a) | 0.602 | --                                                 |

```

8 1.3 (0.106a) 0.702 --
9 1.3 (0.106a) 0.3 22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00

```

| Mod | MAC-Address(es)                        | Serial-Num  |
|-----|----------------------------------------|-------------|
| 2   | 00-05-30-00-9d-d2 to 00-05-30-00-9d-de | JAB064605a2 |
| 5   | 00-05-30-00-64-be to 00-05-30-00-64-c2 |             |
| 6   | 00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd | JAB06350B1R |
| 8   | 00-05-30-01-37-7a to 00-05-30-01-37-fe | JAB072705ja |
| 9   | 00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 | JAB06280ae9 |

\* this terminal session

The `Status` column in the output should display an OK status for switching modules and an active or HA-standby status for supervisor modules. If the status is either OK or active, you can continue with your configuration.

## Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches.

It ensures that the process-level failures do not cause system-level failures. It also restarts the failed processes automatically.

This vital process functions on infrastructure that is internal to the switch.

See [“Displaying System Processes” section on page 31-2](#).

## Synchronizing Supervisor Modules

The running image is automatically synchronized in the standby supervisor module by the active supervisor module. The boot variables are synchronized during this process.

The standby supervisor module, automatically synchronizes its image with the running image on the active supervisor module.



### Note

Refer to the [“Replacing Modules” section on page 6-25](#) for further details.

## Automatically Copying Images to the Standby Supervisor

The **boot auto-copy** copies the boot variable images which are local (present) in the active supervisor module (but not in the standby supervisor module) to the standby supervisor module. Only those KICKSTART and SYSTEM boot variables that are set for the standby supervisor module may be copied. For module (line card) images, all of them that are not present in standby's corresponding locations (bootflash or slot0) will be copied.

To enable or disable automatic copying of boot variables, follow these steps:

|        | Command                                    | Purpose                                                                                                         |
|--------|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)# | Enters configuration mode.                                                                                      |
| Step 2 | switch(config)# <b>boot auto-copy</b>      | Enables automatic copying of boot variables from the active supervisor module to the standby supervisor module. |
|        | switch(config)# <b>no boot auto-copy</b>   | Disables the automatic copy feature (default).                                                                  |

To verify the current state of the auto-copy feature, use the **show boot auto-copy** command (see [Example 5-1](#) and [Example 5-2](#)).

**Example 5-1 Displays the auto-copy Option in an Enabled State**

```
switch# show boot auto-copy
Boot variables Auto-Copy ON
```

**Example 5-2 Displays the auto-copy Option in a Disabled State**

```
switch# show boot auto-copy
Boot variables Auto-Copy OFF
```

To verify what files are being copied, use the **show boot auto-copy list** command. The following example displays image being copied to the standby supervisor module's bootflash, and once this is successful, the next file will be image2.bin. This command only displays files on the active supervisor module (see [Example 5-3](#)).

**Example 5-3 Displays the Files Being Copied**

```
switch# show boot auto-copy list
File: /bootflash:/image1.bin
Bootvar: kickstart

File:/bootflash:/image2.bin
Bootvar: system
```

The following example displays a typical message when the auto-copy option is disabled or if no files are copied (see [Example 5-4](#)).

**Example 5-4 Displays the Current auto-copy State**

```
switch# show boot auto-copy list
No file currently being auto-copied
```

## Displaying HA Information

Use the **show system redundancy status** command to view the high availability status of the system. See [Example 5-5](#). Tables [5-1](#) to [5-3](#) explain the possible redundancy, supervisor, and internal states output in this command.

**Example 5-5 Displays Redundancy Status**

```

switch# show system redundancy status
Redundancy mode

 administrative: HA
 operational: HA
This supervisor (sup-1)

 Redundancy state: Active
 Supervisor state: Active
 Internal state: Active with HA standby
Other supervisor (sup-2)

 Redundancy state: Standby
 Supervisor state: HA standby
 Internal state: HA standby

```

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is *Active with HA standby* and of the other supervisor module is *HA standby*, the switch is operationally HA and can do automatic synchronization.
- If the internal state of one of the supervisor modules is *none* the switch cannot do automatic synchronization.

Table 5-1 lists the possible values for the redundancy states.

**Table 5-1 Redundancy States**

| State        | Description                                                                                                                                                                                            |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Not present  | The supervisor module is not present or is not plugged into the chassis.                                                                                                                               |
| Initializing | The diagnostics have passed and the configuration is being downloaded.                                                                                                                                 |
| Active       | This module is the active supervisor module and the switch is ready to be configured.                                                                                                                  |
| Standby      | This state indicate that a switchover is possible.                                                                                                                                                     |
| Failed       | The switch detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three (3) times. After the third attempt it continues to display a failed state. |
| Offline      | The switch is intentionally shut down for debugging purposes.                                                                                                                                          |
| At BIOS      | The module has established connection with the supervisor and the supervisor module is performing diagnostics.                                                                                         |
| Unknown      | The switch is in an invalid state. If it persists, call TAC.                                                                                                                                           |

Table 5-2 lists the possible values for the Supervisor state.

**Table 5-2 Supervisor States**

| State      | Description                                                                           |
|------------|---------------------------------------------------------------------------------------|
| Active     | This module is the active supervisor module and the switch is ready to be configured. |
| HA standby | This state indicate that a switchover is possible.                                    |



**Table 5-2 Supervisor States (continued)**

| State   | Description                                                           |
|---------|-----------------------------------------------------------------------|
| Offline | The switch is intentionally shut down for debugging purposes.         |
| Unknown | The switch is in an invalid state and requires a support call to TAC. |

[Table 5-3](#) lists the possible values for the internal redundancy state of the supervisor modules.

**Table 5-3 Internal States**

| State                          | Description                                                                                                                                            |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| HA standby                     | This module is the standby supervisor module and the HA switchover mechanism is enabled (see the <a href="#">“HA Switchover”</a> section on page 5-3). |
| Active with no standby         | This state indicate that a switchover is possible.                                                                                                     |
| Active with HA standby         | This module is the active supervisor module and the switch is ready to be configured. The standby module is in the HA-standby state.                   |
| Shutting down                  | The switch is being shut down.                                                                                                                         |
| HA switchover in progress      | The switch is in the process of changing over to the HA switchover mechanism.                                                                          |
| Offline                        | The switch is intentionally shut down for debugging purposes.                                                                                          |
| HA synchronization in progress | The standby supervisor module is in the process of synchronizing its supervisor modules.                                                               |
| Standby (failed)               | The standby supervisor module is not functioning.                                                                                                      |
| Active with failed standby     | This module is the active supervisor module and the second supervisor module is present but is not functioning.                                        |
| Other                          | The switch is in a transient state. If it persists, call TAC.                                                                                          |





## Software Images

---

This chapter describes how to install and upgrade software images. The software image install procedure is dependent on the following factors:

- Software images—The kickstart and system image files reside in directories or folders that can be accessed from the Cisco MDS 9000 Family switch prompt.
- Image version—Each image file has a version.
- Flash disks on the switch—The bootflash: resides on the supervisor and the CompactFlash disk is inserted into the slot0: device.
- Supervisor modules—There are single or dual supervisor modules. In the dual supervisor scenario, the standby supervisor module should be updated first.

This chapter includes the following sections:

- [About Software Images, page 6-2](#)
- [Essential Upgrade Prerequisites, page 6-2](#)
- [Software Upgrade Mechanisms, page 6-4](#)
- [Performing an Automated, One-Step Upgrade, page 6-6](#)
- [Viewing the Status of an Upgrade, page 6-15](#)
- [Performing a Manual Upgrade on a Dual Supervisor Switch, page 6-17](#)
- [Quick, One-Step Upgrade, page 6-23](#)
- [Recovering a Corrupted Bootflash, page 6-26](#)
- [Maintaining Supervisor Modules, page 6-23](#)
- [Replacing Modules, page 6-25](#)
- [Recovery for Switches with Dual Supervisor Modules, page 6-33](#)
- [Default Factory Settings, page 6-35](#)

# About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The SAN-OS consists of two images—the kickstart image and the system image. To upgrade the switch to a new image, you must specify the variables which direct the switch to the images.

- To select the kickstart image use the KICKSTART variable.
- To select the system image use the SYSTEM variable.

The images and variables are important factors in any install procedure. You must specify the variable and the image to upgrade your switch (see the “[Upgrading Modules](#)” section on page 6-22). Both images are not always required for each install.

**Note**

Unless explicitly stated, the software install procedures in this section apply to any switch in the Cisco MDS 9000 Family.

## Essential Upgrade Prerequisites

Before attempting to migrate to any software image version, follow these guidelines:

- Customer Service

Before performing any software upgrade, contact your respective customer service representative to review your software upgrade requirements and to provide recommendations based on your current operating environment.

**Note**

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

- Scheduling

Schedule the upgrade when the fabric is stable and steady. Ensure that everyone who has access to the switch or the network is not configuring the switch or the network during this time. All configurations will be disallowed at this time.

- Space

Verify that sufficient space is available in the location where you are copying the images. This location includes the active and standby supervisor modules or bootflash: (internal to the switch). Use the **dir** command to ensure that the required free space is available for the image files to be copied.

- Standby supervisor module bootflash: directory (see the [Chapter 4, “Initial Configuration”](#)).
- Internal bootflash offers approximately 200 MB of user space.

- Hardware

Avoid power interruption during any install procedure. These kinds of problems can corrupt the software image.

- Connectivity (to retrieve images from remote servers)

- Configure the IP address for the 10/100 BASE-T Ethernet port connection (interface mgmt0).

- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router to route traffic between subnets. Verify connectivity to the remote server using the **ping** command.
- Images
  - The specified system and kickstart images must be compatible with each other.
  - If the kickstart image is not specified, the switch uses the current running kickstart image. If you specify a different system image, ensure that it is compatible with the running kickstart image.
  - Images can be retrieved in one of two ways:
 

Local—images are locally available on the switch (the **install all** command uses the specified local images).

Remote—images are in a remote location and the user specifies the destination using the remote server parameters and the file name to be used locally.
- Commands
  - We recommend the one-step **install all** command to upgrade your software. This command upgrades all modules in any Cisco MDS 9000 Family switch (see the [“The install all command” section on page 6-6](#)).

Only one **install all** command can be running on a switch at any time.

No other command can be issued while running the **install all** command.

The **install all** command cannot be performed on the standby supervisor module—it can only be issued on the active supervisor module.

If the switching module(s) are not compatible with the new supervisor module image, some traffic disruption may be noticed in the related modules, depending on your configuration. These modules are identified in the summary when you issue the **install all** command. You can choose to proceed with the upgrade or abort at this point.

**Note**

When you issue the **install all** command, the switch displays a summary of changes that will be made to your configuration and waits for your authorization to continue executing the command process.

[Table 6-1](#) summarizes terms used in this chapter with specific reference to the install and upgrade process.

**Table 6-1 Terms Specific to this Chapter**

| Term         |           | Definition                                                            |
|--------------|-----------|-----------------------------------------------------------------------|
| bootable     |           | The modules ability to boot or not boot based on image compatibility. |
| impact       |           | The type of software upgrade mechanism—disruptive or non-disruptive.  |
| install-type | reset     | Resets the module.                                                    |
|              | sw-reset  | Resets the module immediately after switchover.                       |
|              | rolling   | Upgrades each module in sequence.                                     |
|              | copy-only | Updates the software for BIOS, loader, or bootrom.                    |

# Software Upgrade Mechanisms

The Cisco MDS SAN-OS software, designed for mission-critical high availability environments, provides the ability to upgrade software without any disruptions. To realize the benefits of nondisruptive upgrades on the Cisco MDS 9509 Director, it is highly recommended that you install dual supervisor modules.

You can upgrade any switch in the Cisco MDS 9500 Series and the Cisco MDS 9200 Series using one of the following mechanisms:

- Automated, one-step upgrade using the **install all** command (see the “[Performing an Automated, One-Step Upgrade](#)” section on page 6-6).
- Quick, one-step upgrade using the **reload** command. This upgrade is disruptive (see the “[Quick, One-Step Upgrade](#)” section on page 6-23).



Tip

The **install all** command compares and presents the results of the compatibility before proceeding with the installation. You have the opportunity to exit, if you do not want to proceed with these changes.

For nondisruptive upgrades, use the automated one-step upgrade. In some cases, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:

- A single supervisor system with kickstart or system image changes
- A dual supervisor system with incompatible system software images.

## Determining Compatibility

If the running image and the image you want to install are incompatible, the **install all** command reports the incompatibility. In some cases, you may decide to proceed with this installation. If the active and the standby supervisor modules run different versions of the image, both images may be HA compatible in some cases and incompatible in others.

Compatibility is established based on the image and configuration:

- Image incompatibility—implies that the running image and the image to be installed are not compatible.
- Configuration incompatibility—implies a possible compatibility if certain features in the running image are turned off as they are not supported in the image to be installed. The image to be installed is considered incompatible with the running image if one of the following statements are true:
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and may cause the switch to move into an inconsistent state. In this case, the incompatibility is *strict*.
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and does not cause the switch to move into an inconsistent state. In this case, the incompatibility is *loose*.

To view the results of a dynamic compatibility check, issue the **show incompatibility system bootflash:filename** command. Use this command to obtain further information when the **install all** command returns the following message:

```
Warning: The startup config contains commands not supported by the standby supervisor; as
a result, some resources might become unavailable after a switchover.
```

```
Do you wish to continue? (y/ n) [y]: n
```

**Example 6-1 Displays HA Compatibility Status**

```
switch# show incompatibility system bootflash:running-image
The following configurations on active are incompatible with the system image
1) Feature Index : 67 , Capability : CAP_FEATURE_SPAN_FC_TUNNEL_CFG
Description : SPAN - Remote SPAN feature using fc-tunnels
Capability requirement : STRICT

2) Feature Index : 119 , Capability : CAP_FEATURE_FC_TUNNEL_CFG
Description : fc-tunnel is enabled
Capability requirement : STRICT
```

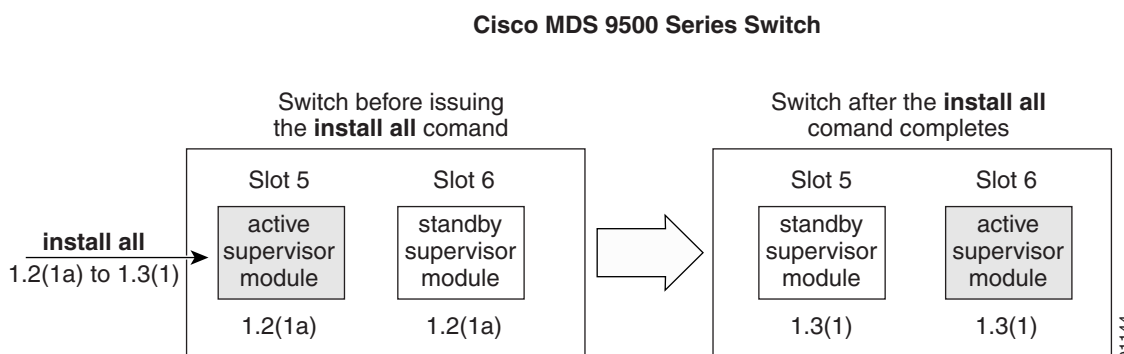
# Performing an Automated, One-Step Upgrade

You can perform an automated upgrade on any switch in the Cisco MDS 9200 Series or the Cisco MDS 9000 Series using the **install all** command.

## The install all command

The **install all** command upgrades all modules in any Cisco MDS 9000 Family switch. [Figure 6-1](#) provides an overview of the switch status before and after issuing the **install all** command.

**Figure 6-1 The install all Command Effect**



The **install all** command automatically verifies if the standby supervisor module is functioning (if present). If it is not functioning, it reloads that module and uses the **reload module slot force-dnld** command to force it to function.



### Note

The **install all** command is only effective on switches running Cisco MDS SAN-OS Release 1.0(3) and later releases.

## Benefits of Using the install all Command

The **install all** command provides the following benefits:

- You can upgrade the entire switch using just one command.
- You will receive descriptive information on the intended changes to your system before you issue the command.
- You have the option to cancel the command. Once the effects of the command are presented, you can choose to continue or cancel when you see this question (the default is **no**):  

```
Do you want to continue y/n? [n] :y
```
- You can upgrade the entire switch using the least disruptive procedure.
- You can view the progress of this command on the console, Telnet, and SSH screens:
  - after a switchover process, you can see the progress from both the supervisor modules.
  - before a switchover process, you can only see the progress from the active supervisor module.



- The image integrity is automatically checked by the **install all** command. This includes the running kickstart and system images.
- A platform validity check is performed to verify that a wrong image is not used—for example, to check if an MDS 9509 image is used to upgrade an MDS 9216 switch.
- Use the **Ctrl-c** escape sequence to gracefully abort the **install all** command. The command sequence completes the update step in progress and returns to the switch prompt. (Other upgrade commands cannot be aborted using **Ctrl-c**.)
- After issuing the **install all** command, if any step in the sequence fails, the command will complete the step in progress and abort.

For example, if a switching module fails to be updated for any reason (for example, due to an unstable fabric state), then the command sequence disruptively updates that module and aborts. In such cases, you can verify the problem on the affected switching module and manually upgrade the other switching modules.

## Recognizing Failure Cases

The following situations will cause the **install all** command to abort:

- If the standby supervisor module bootflash: directory does not have sufficient space to accept the updated image.
- If the specified system and kickstart images are not compatible.
- If the **install all** command is issued on the standby supervisor module.
- If the fabric or switch is configured while the upgrade is in progress.
- If a module is removed while the upgrade is in progress.
- If the switch has any power disruption while the upgrade is in progress.
- If the entire path for the remote location is not specified accurately.
- If images are incompatible after an upgrade. For example, a switching module image may be incompatible with the system image, or a kickstart image may be incompatible with a system image. This is also identified by the **show install all impact** command in the compatibility check section of the output (under the Bootable column).



### Caution

Avoid aborting the switch progress after issuing the **install all** command. If the **install all** command is aborted, be sure to verify the state of the switch at every stage and reissue the command after 10 seconds. If you reissue the **install all** command within the 10-second span, the command is rejected with an error message indicating that an installation is currently in progress.

## Using the install all Command

To perform an automated software upgrade on any switch, follow these steps:

- 
- Step 1** Log into the switch through the console port of the active supervisor.
  - Step 2** Create a backup of your existing configuration file, if required (see the [“Working with Configuration Files”](#) section on page 4-23).
  - Step 3** Perform the upgrade by issuing the **install all** command.

```

switch# install all sys bootflash:isan-1.3.1 kickstart bootflash:boot-1.3.1

Verifying image bootflash:/boot-1.3.1
[#####] 100% -- SUCCESS

Verifying image bootflash:/isan-1.3.1
[#####] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS

Extracting "ips" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS

Extracting "loader" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS

Compatibility check is done:
Module bootable Impact Install-type Reason
----- -
 1 yes non-disruptive rolling
 2 yes disruptive rolling Hitless upgrade is not supported
 3 yes disruptive rolling Hitless upgrade is not supported
 4 yes non-disruptive rolling
 5 yes non-disruptive reset
 6 yes non-disruptive reset

Images will be upgraded according to following table:
Module Image Running-Version New-Version Upg-Required
----- -
 1 slc 1.3(2a) 1.3(1) yes
 1 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
 2 ips 1.3(2a) 1.3(1) yes
 2 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
 3 ips 1.3(2a) 1.3(1) yes
 3 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
 4 slc 1.3(2a) 1.3(1) yes
 4 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
 5 system 1.3(2a) 1.3(1) yes
 5 kickstart 1.3(2a) 1.3(1) yes
 5 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
 5 loader 1.2(2) 1.2(2) no
 6 system 1.3(2a) 1.3(1) yes
 6 kickstart 1.3(2a) 1.3(1) yes
 6 bios v1.1.0(10/24/03) v1.1.0(10/24/03) no
 6 loader 1.2(2) 1.2(2) no

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Syncing image bootflash:/boot-1.3.1 to standby.
[#####] 100% -- SUCCESS

Syncing image bootflash:/isan-1.3.1 to standby.
[#####] 100% -- SUCCESS
Jan 18 23:40:03 Hacienda %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from

```

```
Performing configuration copy.
[#####] 100% -- SUCCESS

Module 6: Waiting for module online.
|
Auto booting bootflash:/boot-1.3.1 bootflash:/isan-1.3.1...
Booting kickstart image: bootflash:/boot-1.3.1....
.....Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..r.r.. done.
Loading system software
Uncompressing system image: bootflash:/isan-1.3.1
ccc
cccccccccccccccccccccccccc
INIT: Entering runlevel: 3
```

**Step 4** Exit the switch console and open a new terminal session to view the upgraded supervisor module using the **show module** command.



### Note

If the configuration meets all guidelines when the **install all** command is issued, all modules (supervisor and switching) are upgraded. This is true for any switch in the Cisco MDS 9000 Family.



### Caution

If a nondisruptive upgrade operation failed for any reason other than those listed in the [“Recognizing Failure Cases”](#) section on page 6-7, contact your reseller or Cisco representative for further assistance.

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL:  
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>



### Note

Any software upgrade for the Caching Services Module (CSM) and the IP Storage (IPS) services module is disruptive. These modules use a rolling upgrade install mechanism whereby, both modules can only be upgraded in sequence. After the first module is upgraded and before the second module is going to be upgraded, the SAN-OS software introduces a time delay to ensure that all applications in the module reach a steady state. The IPS module require a 5-minute delay before the next IPS module is upgraded to guarantee a stable state.

The CSM module requires a 30 minute delay before the next CSM module is upgraded to guarantee a stable state. Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for further information on upgrading CSMs.

## Sample install all Commands

[Example 6-2](#) displays the result of the **install all** command issued from a console terminal that is connected to the active supervisor. Once a switchover happens, you can see the rest of the output from the console terminal of the standby supervisor module. [Example 6-3](#) displays the file output continuation of the **install all** command on the console of the standby supervisor module.

Similarly, you can view the results of the **install all** command issued from the SSH or Telnet terminal that is connected to the active supervisor. Once a switchover happens, you need to log back into the switch and issue the **show install all status** command (see the [Viewing the Status of an Upgrade](#), page 6-15).

### Example 6-2 Successful install all Command Issued from the Active Console

```
Hacienda# install all sys bootflash:isan-1.3.1 kickstart bootflash:boot-1.3.1

Verifying image bootflash:/boot-1.3.1
[#####] 100% -- SUCCESS

Verifying image bootflash:/isan-1.3.1
[#####] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS

Extracting "ips" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS

Extracting "loader" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS

Compatibility check is done:
Module bootable Impact Install-type Reason
----- -
1 yes non-disruptive rolling
2 yes disruptive rolling Hitless upgrade is not supported
3 yes disruptive rolling Hitless upgrade is not supported
4 yes non-disruptive rolling
5 yes non-disruptive reset
6 yes non-disruptive reset

Images will be upgraded according to following table:
Module Image Running-Version New-Version Upg-Required
----- -
1 slc 1.3 (2a) 1.3 (1) yes
1 bios v1.1.0 (10/24/03) v1.1.0 (10/24/03) no
2 ips 1.3 (2a) 1.3 (1) yes
2 bios v1.1.0 (10/24/03) v1.1.0 (10/24/03) no
3 ips 1.3 (2a) 1.3 (1) yes
3 bios v1.1.0 (10/24/03) v1.1.0 (10/24/03) no
4 slc 1.3 (2a) 1.3 (1) yes
```



```

Module 6: Waiting for module online.
Jan 18 23:43:02 Hacienda %PORT-5-IF_UP: Interface mgmt0 is up
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature
FM_SERVER_PKG. Application(s) shutdown in 53 days.
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature
ENTERPRISE_PKG. Application(s) shutdown in 50 days.
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature
SAN_EXTN_OVER_IP. Application(s) shutdown in 50 days.
Jan 18 23:43:19 Hacienda %LICMGR-3-LOG_LICAPP_NO_LIC: Application port-security running
without ENTERPRISE_PKG license, shutdown in 50 days
Jan 18 23:43:19 Hacienda %LICMGR-4-LOG_LICAPP_EXPIRY_WARNING: Application Roles evaluation
license ENTERPRISE_PKG expiry in 50 days
Jan 18 23:44:54 Hacienda %BOOTVAR-5-NEIGHBOR_UPDATE_AUTOCOPY: auto-copy supported by
neighbor, starting...
Module 1: Non-disruptive upgrading.
[#] 0%Jan 18 23:44:56 Hacienda %MODULE-5-STANDBY_SUP_OK: Supervisor 5
is standby
Jan 18 23:44:55 Hacienda %IMAGE_DNLD-SLOT1-2-IMG_DNLD_STARTED: Module image download
process. Please wait until completion...
Jan 18 23:45:12 Hacienda %IMAGE_DNLD-SLOT1-2-IMG_DNLD_COMPLETE: Module image download
process. Download successful.
Jan 18 23:45:48 Hacienda %MODULE-5-MOD_OK: Module 1 is online
[#####] 100% -- SUCCESS
Module 4: Non-disruptive upgrading.
[#] 0%Jan 18 23:46:12 Hacienda %IMAGE_DNLD-SLOT4-2-IMG_DNLD_STARTED:
Module image download process. Please wait until completion...
Jan 18 23:46:26 Hacienda %IMAGE_DNLD-SLOT4-2-IMG_DNLD_COMPLETE: Module image download
process. Download successful.
Jan 18 23:47:02 Hacienda %MODULE-5-MOD_OK: Module 4 is online
[#####] 100% -- SUCCESS
Module 2: Disruptive upgrading.
...
-- SUCCESS
Module 3: Disruptive upgrading.
...
-- SUCCESS
Install has been successful.
MDS Switch
Hacienda login:

```

**Example 6-4** displays the result of the **install all** command if the system and kickstart files are automatically downloaded using a remote (TFTP, FTP, SCP, or SFTP) download option.



#### Caution

Specify the exact, complete, path of the remote location. The system will not allow you to proceed if the entire path is not accurately specified. Some examples of incomplete **install all** commands are provided below, [Example 6-4](#) provides an accurate, complete example.

```

switch# install all system bootflash:system-image kickstart tftp:
Please provide a complete URI
switch# install all system scp:
Please provide a complete URI

```

#### **Example 6-4 A Sample of the install all Command Issued Using a Remote Download**

```

switch# install all system
scp://user@171.69.16.26/tftpboot/HKrel/qa/vegas/final/m9500-sflek9-mz.1.3.2a.bin kickstart
scp://user@171.69.16.26/tftpboot/HKrel/qa/vegas/final/m9500-sflek9-kickstart-mz.1.3.2a.bin
For scp://user@171.69.16.26, please enter password:
For scp://user@171.69.16.26, please enter password:

```

```

Copying image from
scp://user@171.69.16.26/tftpboot/HKrel/qa/vegas/final/m9500-sflek9-kickstart-mz.1.3.2a.bin
to bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Copying image from
scp://user@171.69.16.26/tftpboot/HKrel/qa/vegas/final/m9500-sflek9-mz.1.3.2a.bin to
bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Verifying image bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin
[#####] 100% -- SUCCESS

Verifying image bootflash:///m9500-sflek9-mz.1.3.2a.bin
[#####] 100% -- SUCCESS

Extracting "slc" version from image bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "ips" version from image bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image
bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

```

Compatibility check is done:

| Module | bootable | Impact         | Install-type | Reason                           |
|--------|----------|----------------|--------------|----------------------------------|
| -----  | -----    | -----          | -----        | -----                            |
| 1      | yes      | non-disruptive | rolling      |                                  |
| 2      | yes      | disruptive     | rolling      | Hitless upgrade is not supported |
| 3      | yes      | non-disruptive | rolling      |                                  |
| 4      | yes      | non-disruptive | rolling      |                                  |
| 5      | yes      | non-disruptive | reset        |                                  |
| 6      | yes      | non-disruptive | reset        |                                  |
| 7      | yes      | non-disruptive | rolling      |                                  |
| 8      | yes      | non-disruptive | rolling      |                                  |
| 9      | yes      | disruptive     | rolling      | Hitless upgrade is not supported |

Images will be upgraded according to following table:

| Module | Image     | Running-Version   | New-Version       | Upg-Required |
|--------|-----------|-------------------|-------------------|--------------|
| -----  | -----     | -----             | -----             | -----        |
| 1      | slc       | 1.3 (1)           | 1.3 (2a)          | yes          |
| 1      | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no           |
| 2      | ips       | 1.3 (1)           | 1.3 (2a)          | yes          |
| 2      | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no           |
| 3      | slc       | 1.3 (1)           | 1.3 (2a)          | yes          |
| 3      | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no           |
| 4      | slc       | 1.3 (1)           | 1.3 (2a)          | yes          |
| 4      | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no           |
| 5      | system    | 1.3 (1)           | 1.3 (2a)          | yes          |
| 5      | kickstart | 1.3 (1)           | 1.3 (2a)          | yes          |
| 5      | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no           |
| 5      | loader    | 1.2 (2)           | 1.2 (2)           | no           |

|   |           |                   |                   |     |
|---|-----------|-------------------|-------------------|-----|
| 6 | system    | 1.3 (1)           | 1.3 (2a)          | yes |
| 6 | kickstart | 1.3 (1)           | 1.3 (2a)          | yes |
| 6 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 6 | loader    | 1.2 (2)           | 1.2 (2)           | no  |
| 7 | slc       | 1.3 (1)           | 1.3 (2a)          | yes |
| 7 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 8 | slc       | 1.3 (1)           | 1.3 (2a)          | yes |
| 8 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 9 | ips       | 1.3 (1)           | 1.3 (2a)          | yes |
| 9 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |

Do you want to continue with the installation (y/n)? [n]

**Example 6-5** displays the **install all** command output of a failed operation due to a lack of disk space.

#### **Example 6-5 Failed Operation due to a Full bootflash: Directory**

```
switch# install all system bootflash:isan-1.3.2a kickstart bootflash:boot-1.3.2a
```

```
Verifying image bootflash:/boot-1.3.2a
[#####] 100% -- SUCCESS
```

```
Verifying image bootflash:/isan-1.3.2a
[#####] 100% -- SUCCESS
```

```
Extracting "slc" version from image bootflash:/isan-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "ips" version from image bootflash:/isan-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:/isan-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "kickstart" version from image bootflash:/boot-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "loader" version from image bootflash:/boot-1.3.2a.
[#####] 100% -- SUCCESS
```

Compatibility check is done:

| Module | bootable | Impact         | Install-type | Reason                           |
|--------|----------|----------------|--------------|----------------------------------|
| 1      | yes      | non-disruptive | rolling      |                                  |
| 2      | yes      | disruptive     | rolling      | Hitless upgrade is not supported |
| 3      | yes      | non-disruptive | rolling      |                                  |
| 4      | yes      | non-disruptive | rolling      |                                  |
| 5      | yes      | non-disruptive | reset        |                                  |
| 6      | yes      | non-disruptive | reset        |                                  |
| 7      | yes      | non-disruptive | rolling      |                                  |
| 8      | yes      | non-disruptive | rolling      |                                  |
| 9      | yes      | disruptive     | rolling      | Hitless upgrade is not supported |

Images will be upgraded according to following table:

| Module | Image | Running-Version   | New-Version       | Upg-Required |
|--------|-------|-------------------|-------------------|--------------|
| 1      | slc   | 1.3 (1)           | 1.3 (2a)          | yes          |
| 1      | bios  | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no           |
| 2      | ips   | 1.3 (1)           | 1.3 (2a)          | yes          |



|   |           |                   |                   |     |
|---|-----------|-------------------|-------------------|-----|
| 2 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 3 | slc       | 1.3 (1)           | 1.3 (2a)          | yes |
| 3 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 4 | slc       | 1.3 (1)           | 1.3 (2a)          | yes |
| 4 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 5 | system    | 1.3 (1)           | 1.3 (2a)          | yes |
| 5 | kickstart | 1.3 (1)           | 1.3 (2a)          | yes |
| 5 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 5 | loader    | 1.2 (2)           | 1.2 (2)           | no  |
| 6 | system    | 1.3 (1)           | 1.3 (2a)          | yes |
| 6 | kickstart | 1.3 (1)           | 1.3 (2a)          | yes |
| 6 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 6 | loader    | 1.2 (2)           | 1.2 (2)           | no  |
| 7 | slc       | 1.3 (1)           | 1.3 (2a)          | yes |
| 7 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 8 | slc       | 1.3 (1)           | 1.3 (2a)          | yes |
| 8 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |
| 9 | ips       | 1.3 (1)           | 1.3 (2a)          | yes |
| 9 | bios      | v1.1.0 (10/24/03) | v1.0.8 (08/07/03) | no  |

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Syncing image bootflash:/boot-1.3.2a to standby.

[#####] 100% -- SUCCESS

Syncing image bootflash:/isan-1.3.2a to standby.

[# ] 0% -- FAIL. Return code 0x401E0008 (request was aborted, standby disk may be full).

Install has failed. Return code 0x40930013 (Syncing images to standby failed). <----

#### -----insufficient space message

Please identify the cause of the failure, and try 'install all' again.

Dec 15 19:36:42 switch %SYSMGR-3-SERVICE\_TERMINATED: Service "installer" (PID 5470) has finished with error code SYSMGR\_EXITCODE\_FAILURE\_NOCALLHOME (20).

**Example 6-6** displays the **install all** command output of a failed operation due to an invalid image.

#### **Example 6-6 Failed Operation due to an Invalid Image**

```
install all sys bootflash:junk kickstart bootflash:junk
```

Verifying image bootflash:/junk

[# ] 0% -- FAIL. Return code 0x4045001E (mismatch between actual image type and boot variable).

Compatibility check failed. Return code 0x40930011 (Image verification failed).

Hacienda# Jan 19 00:20:35 Hacienda %SYSMGR-3-SERVICE\_TERMINATED: Service "installer" (PID 5664) has finished with error code SYSMGR\_EXITCODE\_FAILURE\_NOCALLHOME (20).

## Viewing the Status of an Upgrade

To view the on-going **install all** command or the log of the last installed **install all** command from a Console, SSH, or Telnet session, use the **show install all status** command.

This command allows you to view the **install all** output on both the active and standby supervisor module even if you are not connected to the console terminal. It only displays the status of an **install all** command that is issued from the CLI (not the GUI).

**Example 6-7 Displays the install all Command Output**

```
switch# show install all status
There is an on-going installation... <----- in progress installation
Enter Ctrl-C to go back to the prompt.

Verifying image bootflash:/b-1.3.0.104
-- SUCCESS

Verifying image bootflash:/i-1.3.0.104
-- SUCCESS

Extracting "system" version from image bootflash:/i-1.3.0.104.
-- SUCCESS

Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
-- SUCCESS

Extracting "loader" version from image bootflash:/b-1.3.0.104.
-- SUCCESS

switch# show install all status
This is the log of last installation. <<<<< log of last install

Verifying image bootflash:/b-1.3.0.104
-- SUCCESS

Verifying image bootflash:/i-1.3.0.104
-- SUCCESS

Extracting "system" version from image bootflash:/i-1.3.0.104.
-- SUCCESS

Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
-- SUCCESS

Extracting "loader" version from image bootflash:/b-1.3.0.104.
-- SUCCESS
```

# Performing a Manual Upgrade on a Dual Supervisor Switch



## Caution

If you are a new user, use the **install all** command to perform a software upgrade. The information and instructions provided in this section targets administrators or individuals who are completely familiar with specific switch functions.

You can manually upgrade any switch in the Cisco MDS 9200 Series or the Cisco MDS 9500 Series using the procedures provided in this section. This upgrade mechanism requires you to implement some or all procedures depending on your switch or network configuration.

To perform a manual upgrade for a dual supervisor switch, follow these steps.

- 
- Step 1**    [“Preparing for a Manual Installation” section on page 6-17](#)
  - Step 2**    [“Upgrading a Loader” section on page 6-19](#)
  - Step 3**    [“Upgrading the BIOS” section on page 6-21](#)
  - Step 4**    [“Upgrading Modules” section on page 6-22](#)
- 

## Preparing for a Manual Installation

To prepare any Cisco MDS 9000 Family switch for a manual software installation, follow these steps:

- 
- Step 1**    Log into the switch through the console port, an SSH session, or a Telnet session.
  - Step 2**    Create a backup of your existing configuration file, if required (see the [“Saving the Configuration” section on page 4-26](#)).
  - Step 3**    Copy the software image from a TFTP location to one of two targets: bootflash: or slot0:.  
The switch remains operational while the image file is copied.

- Bootflash device (TFTP defaults to the bootflash device)—Copy the software image file from the appropriate TFTP directory to bootflash.

```
switch# copy scp://<server IP address>/<file name in TFTP> <destination file name as desired>
```

For example:

```
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
```



**Note** The Cisco MDS 9216 Switch does not have an external CompactFlash (see the [“Working with Configuration Files” section on page 4-23](#)). If you are using a switch in this series, use the bootflash: directory to copy and verify files.

- CompactFlash device—Copy the software image file from the appropriate TFTP directory to the CompactFlash device in slot0:.

```
switch# copy scp://<server IP address>/<file name in SCP> slot0:system-image
```

You can also copy the image onto a new Flash disk from a PC and insert it in slot0: in the Cisco MDS 9500 Series switch. After you copy the image and insert it into the slot0: partition, the process is the same as the CompactFlash device after the copy command is issued.

**Step 4** Verify that the file was copied in the required directory.

```
switch# dir bootflash:
40295206 Aug 05 15:23:51 1980 ilc1.bin
12456448 Jul 30 23:05:28 1980 kickstart-image1
12288 Jun 23 14:58:44 1980 lost+found/
27602159 Jul 30 23:05:16 1980 system-image1
12447232 Aug 05 15:08:30 1980 kickstart-image2
28364853 Aug 05 15:11:57 1980 system-image2

Usage for bootflash://sup-local
 135404544 bytes used
 49155072 bytes free
184559616 bytes total
```

**Step 5** Ensure that the software images are not damaged or corrupted in the saved bootflash: location.

When copying a new image to your switch, confirm that the image was not corrupted during the copy process.

Use the **show version image** command to verify that the required image was copied successfully.

```
switch# show version image bootflash:kickstart-image
image name: m9500-sflek9-kickstart-mzg.1.0.3.bin
kickstart: version 1.0(3)
loader: version 1.0(3)
compiled: 2/12/2003 11:00:00
```



**Note** A verification failed message is generated when you use a Cisco MDS 9500 Series image on a Cisco MDS 9200 Series switch or a Cisco MDS 9200 Series image on a Cisco MDS 9500 Series switch. Be sure to verify the right image.

**Step 6** Compare the running system image and the new image by issuing the **show install all impact** command.

```
switch# show install all impact
```

## Upgrading a Loader

Using the **install module slot# of the supervisor module loader** command, you can upgrade the (boot) loader.



### Note

If the loader is upgraded, you need to reboot to make the new loader effective. You can schedule the reboot at a convenient time so traffic will not be impacted.



### Caution

Before issuing this command, be sure to read the release notes to verify compatibility issues between the boot loader and the kickstart or system images.

To upgrade the loader on either the active or standby supervisor module, follow these steps.

- Step 1** Use the **show version** command to verify the version on the active and standby supervisor modules.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
 BIOS: version 1.0.8
 loader: version 1.1(2) <-----current running version
 kickstart: version 2.0(1)
 system: version 2.0(1)

 BIOS compile time: 08/07/03
 kickstart image file is: bootflash:///m9500-sflek9-kickstart-mzg.2.0.0.6.bin
 kickstart compile time: 10/25/2010 12:00:00
 system image file is: bootflash:///m9500-sflek9-mzg.2.0.0.6.bin
 system compile time: 10/25/2020 12:00:00

Hardware
 RAM 1024584 kB

 bootflash: 1000944 blocks (block size 512b)
 slot0: 0 blocks (block size 512b)

 172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

 Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
 Reason: Reset Requested by CLI command reload
 System version: 2.0(0.6)
 Service:
```

- Step 2** Issue the **install module** command for the required supervisor module (active or standby). This example displays the command being issued for the standby supervisor module in slot 6.

```
switch# install module 6 loader bootflash:kickstart-image
```

**Note**

If you install a loader version that is the same as the currently-installed version, the command will not execute. When both the current version and the installed version are the same, use the **init system** command to force a loader upgrade.

**Step 3** Use the **show version** command to verify the updated image on the supervisor module.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
 BIOS: version 1.3.1
 loader: version 1.3(1) <-----New running version
 kickstart: version 1.3(1)]
 system: version 1.3(1)

 BIOS compile time: 08/07/03
 kickstart image file is: bootflash:///m9500-sflek9-kickstart-mzg.2.0.0.6.bin
 kickstart compile time: 10/25/2010 12:00:00
 system image file is: bootflash:///m9500-sflek9-mzg.2.0.0.6.bin
 system compile time: 10/25/2020 12:00:00

Hardware
 RAM 1024584 kB

 bootflash: 1000944 blocks (block size 512b)
 slot0: 0 blocks (block size 512b)

172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
Reason: Reset Requested by CLI command reload
System version: 2.0(0.6)
Service:
```

## Upgrading the BIOS



### Tip

Refer to the Release Notes to verify if the BIOS has changed for the image version being used.

Program the supervisor or switching module BIOS only if a new BIOS image is provided by Cisco. Only use the provided image to upgrade the BIOS. This command does not affect traffic and can be issued at any time on any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.



### Note

If the BIOS is upgraded, reboot to make the new BIOS effective. You can schedule the reboot at a convenient time so traffic will not be impacted.



### Caution

The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

To upgrade the BIOS for a module, follow these steps:

**Step 1** Use the **show version** command to verify the current running BIOS version.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.
Software
 BIOS: version 1.0(6) <----- current running version
 loader: version 1.0(3)
 kickstart: version 1.0(3)
 system: version 1.0(3)

 BIOS compile time: 01/27/03
 kickstart image file is: bootflash:/kickstart-image
 kickstart compile time: 01/25/2003 12:00:00
 system image file is: bootflash:/system-image
 system compile time: 01/25/2003 12:00:00

Hardware
 RAM 1027564 kB
```

**Step 2** Verify that the BIOS version of the system image is different from the running image.

```
switch# show version image bootflash:system-image
 image name: m9500-sflek9-mz.1.0.3.bin
 bios: version v1.0.6(01/27/03) <----- BIOS is same version 1.0.6
 system: version 1.0(3)
 compiled: 2/28/2003 5:00:00

system service's list

package name package version
acl 1.0(3)
ascii-cfg 1.0(3)
bios_daemon 1.0(3)
...
```

**Note**

If the version are different, issue the **install module** command as specified in step 3. If they are the same, you do not need to update the BIOS image.

- Step 3** Run the **install module slot# bios** command to install each module (if required). In this example, the supervisor module in slot 6 was updated.

```
switch# install module 6 bios system bootflash:system-image
Started bios programming please wait
[#####] 100%
BIOS upgrade succeeded for module 1
```

**Caution**

Do not reboot the switch if any errors were indicated in response to this command.

- Step 4** Issue the **show version** command to verify that the module was updated with a the new BIOS version.

```
switch# show version module 6
ModNo Image Type SW Version SW Interim Version BIOS Version
6 Stby Sup 1.3(2) 1.3(1.1) 1.1.0 [last 1.0.6]
```

## Upgrading Modules

The **install module slot-number image** command upgrades the switching or services module image to be the same as the supervisor module image—if they are not the same.

To upgrade modules, follow these steps:

- Step 1** Use the **show module** command to determine if a switching module must be updated.
- If the switching module is running an image version that is older than the image version on the supervisor module, that switching module must be updated using this procedure.
- Step 2** Issue the **install module slot# image** command on the module (line card). This command updates the specified switching module using the image version that is running on the active supervisor module.

```
switch# install module 8 image
SW version of linecard image in Supervisor = 1.0(3)
Module 8 : SW version = 1.0(3) , Image version is same, upgrade is not required
```

The preceding example displays a module that does not need to be updated. The following example displays a module that requires an update.

```
switch# install module 9 image
SW version of linecard image in Supervisor = 1.1(1)
Module 9 : SW version = 1.0(3) , Non disruptive image upgrade is possible
Starting non disruptive upgrade of module 9
Feb 3 02:37:13 switch %LC-2-MSG:SLOT9 LC-2-IMG_DNLD_STARTED: Module image
download process. Please wait until completion...
Feb 3 02:37:33 switch %LC-2-MSG:SLOT9 LC-2-IMG_DNLD_COMPLETE: Module image
download process. Download successful.
Module 9 upgrade successful
```

Repeat this step for each switching module that must be updated.



- Step 3** Issue the **show module** to verify the module version command to ensure the modules are functioning as required.
- 

## Quick, One-Step Upgrade

To perform a quick, one-step upgrade on a Cisco MDS 9000 Family switch, follow these steps:

- Step 1** Copy the kickstart and system image files to the required location (see the [“Copying Files” section on page 4-27](#)).
- Step 2** Set the boot variables (see the [“Upgrading Modules” section on page 6-22](#)).
- Step 3** Issue the **reload** command. The **reload** command reboots the system. This upgrade is disruptive.



**Tip** Use the **install all** command to gracefully reload the switch and handle configuration conversions.

---

## Maintaining Supervisor Modules

This section includes general information about replacing and using supervisor modules effectively.

### Standby Supervisor’s Boot Variable Version

If the standby supervisor module’s boot variable images are not the *same* version as those running on the active supervisor module, the software forces the standby supervisor module to run the same version as the active supervisor module.

If you specifically set the boot variables of the standby supervisor module to a different version and reboot the standby supervisor module, the standby supervisor module will only load the specified boot variable if the same version is also running on the active supervisor module. At this point, the standby supervisor module is *not* running the images set in the boot variables.

### Standby Supervisor Boot Alert

If a standby supervisor module fails to boot, the active supervisor module detects that conditions and generates a Call Home event and a SYSLOG message and reboots the standby supervisor module approximately 3 to 6 minutes after the standby supervisor module moves to the `loader>` prompt.

The following SYSLOG error message is issued:

```
%DAEMON-2-SYSTEM_MSG:Standby supervisor failed to boot up.
```

This error message is also generated if one of the following situations apply:

- you remain at the `loader>` prompt for an extended period of time, or

- you do not set the boot variables appropriately.

# Replacing Modules

When you replace any module (supervisor, switching, or services module), you must ensure that the new module is running the same software version as the rest of the switch.

Refer to *Cisco MDS 9000 Family San Volume Controller Configuration Guide* for configuration details on replacing the Caching Services Module (CSM).

**Note**

When a spare standby supervisor module is inserted, it uses the same image as the active supervisor module. The SAN-OS software image is not copied to the standby flash until you issue an **install all** command.

Issuing the **install all** command after replacing any module, ensures the following actions:

- The proper system and kickstart images are copied on the standby bootflash:.
- The proper boot variables are set.
- The loader and the BIOS are upgraded to the same version available on the active supervisor module.

To replace any module in any Cisco MDS 9216 switch or any director in the Cisco MDS 9500 Series, follow these steps:

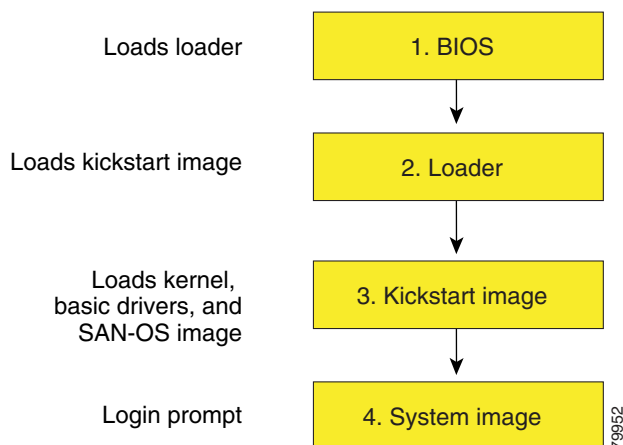
- 
- |               |                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Create a backup of your existing configuration file, if required using the <b>copy running-config startup-config</b> command.                  |
| <b>Step 2</b> | Replace the required module as specified in the <i>Cisco MDS 9216 Switch</i> or the <i>Cisco MDS 9500 Series Hardware Installation Guide</i> . |
| <b>Step 3</b> | Issue the <b>install all</b> command to ensure that the new module is running the same software as the rest of the switch.                     |
| <b>Step 4</b> | Wait till the new module is online and then ensure that the replacement was successful using the <b>show module</b> command.                   |
-

# Recovering a Corrupted Bootflash

All switch configurations reside in the internal bootflash. If you have a corrupted internal bootflash, you could potentially lose your configuration. Be sure to save and back up your configuration files periodically. The regular switch boot goes through the following sequence (see [Figure 6-2](#)):

1. The basic input/output system (BIOS) loads the loader.
2. The boot loader loads the kickstart image into RAM and starts the kickstart image.
3. The kickstart image loads and starts the system image.
4. The system image reads the startup configuration file.

**Figure 6-2 Regular Boot Sequence**



If the images on your switch are corrupted and you are not able to proceed (error state), you can determine the reason and attempt to interrupt the switch boot sequence and recover the image by entering the BIOS configuration utility described in the following section. Access this utility only when needed to recover a corrupted internal disk.

**Caution**

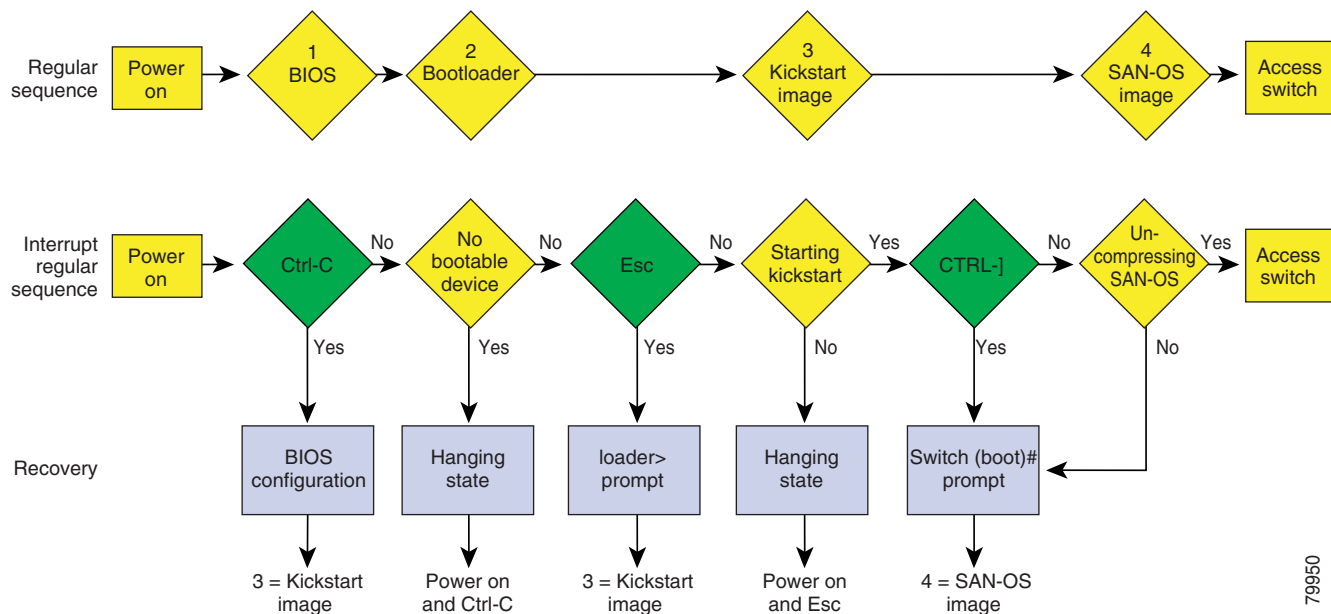
The BIOS changes explained in this section are only required to recover a corrupted bootflash.

Recovery procedures require the regular sequence to be interrupted. The internal switch sequence goes through four phases between the time you turn the switch on and the time the switch prompt appears on your terminal—BIOS, boot loader, Kickstart, and system (see [Table 6-2](#) and [Figure 6-3](#)).

**Table 6-2** Recovery Interruption

| Phase       | Normal Prompt <sup>1</sup> | Recovery Prompt <sup>2</sup> | Description                                                                                                                                                                                                                                                           |
|-------------|----------------------------|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BIOS        | loader>                    | No bootable device           | The BIOS begins the power-on self test, memory test, and other operating system applications. While the test is in progress, press <b>Ctrl-C</b> to enter the BIOS configuration utility and use the netboot option.                                                  |
| Boot loader | Starting kickstart         | loader>                      | The boot loader uncompresses loaded software to boot an image using its file name as reference. These images are made available through bootflash. When the memory test is over, press <b>Esc</b> to enter the boot loader prompt.                                    |
| Kickstart   | Uncompressing system       | switch (boot) #              | When the boot loader phase is over, press <b>Ctrl-]</b> <sup>3</sup> (Control key plus right bracket key) to enter the <code>switch (boot) #</code> prompt. If the corruption causes the console to stop at this prompt, copy the system image and reboot the switch. |
| System      | Login:                     | —                            | The system image loads the configuration file of the last saved running configuration and returns a switch login prompt.                                                                                                                                              |

1. This prompt or message appears at the end of each phase.
2. This prompt or message appears when the switch cannot progress to the next phase.
3. Depending on your telnet client, these keys may differ. Refer to the documentation provided by your telnet client.

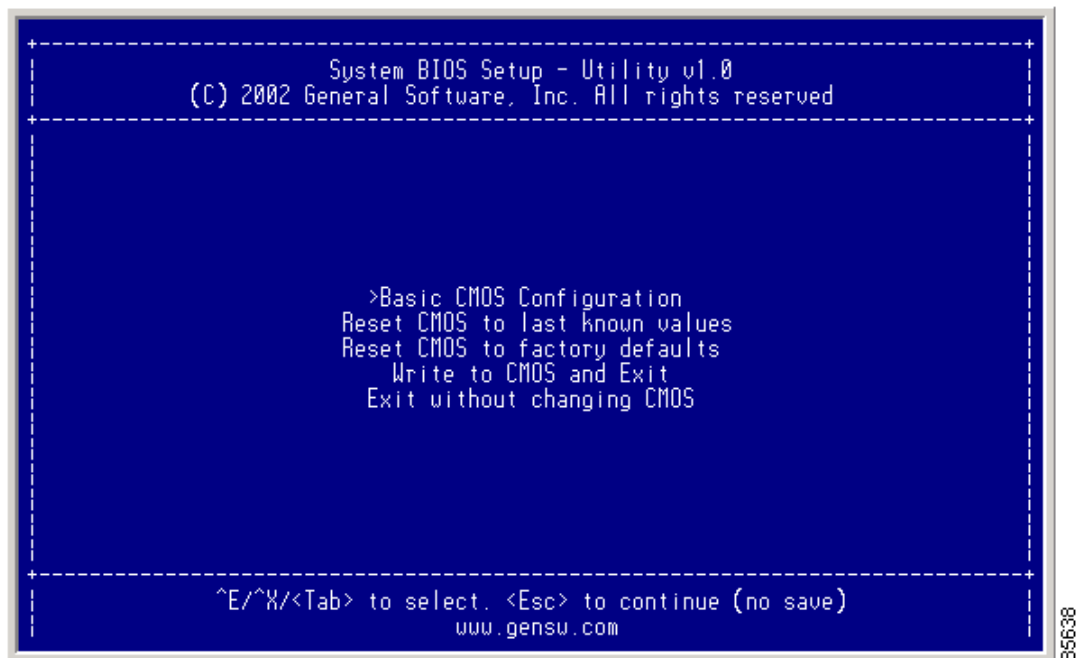
**Figure 6-3** Regular and Recovery Sequence

## Recovery Using BIOS Setup

To recover a corrupted bootflash image (no bootable device found message) for a switch with a single supervisor module, follow these steps:

- Step 1** Connect to the console port of the required switch.
- Step 2** Boot or reboot the switch.
- Step 3** Press **Ctrl-C** to interrupt the BIOS setup during the BIOS memory test.
- You see the netboot BIOS Setup Utility screen (see [Figure 6-4](#)).

**Figure 6-4 BIOS Setup Utility**



**Note**

Your navigating options are provided at the bottom of the screen.

Tab = Jump to next field

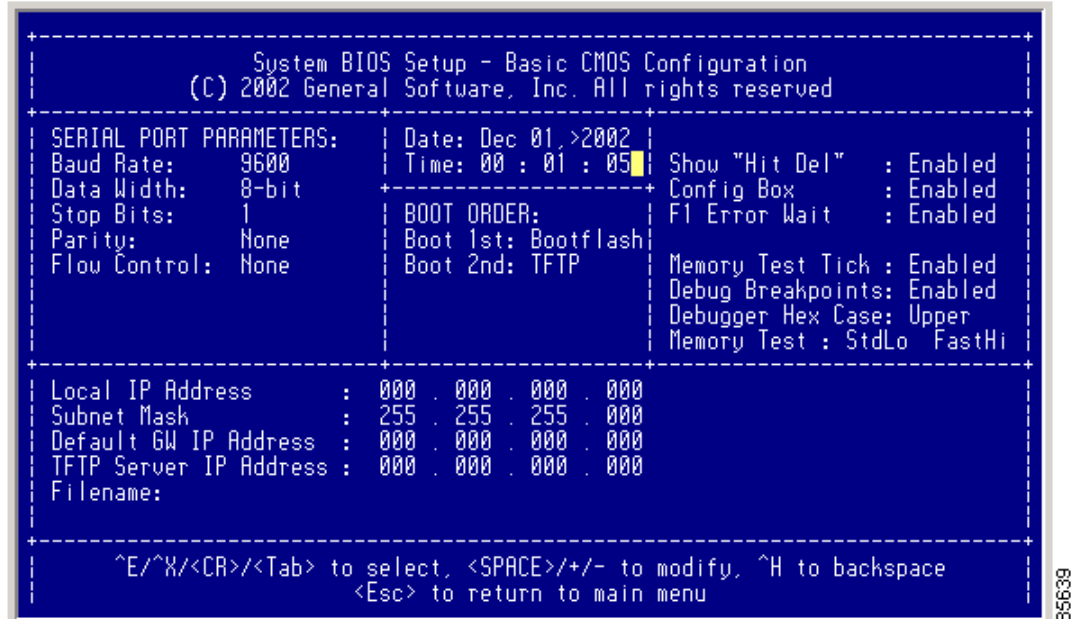
Ctrl-E = Down arrow

Ctrl-X = Up arrow

Ctrl-H = Erase (Backspace might not work if your terminal is not configured properly.)

- Step 4** Press the **Tab** key to select the Basic CMOS Configuration, and press **Enter**.  
You see the BIOS setup CMOS Configuration screen (see Figure 6-5).

**Figure 6-5 BIOS Setup Configuration (CMOS)**



- Step 5** Change the “Boot 1st:” field to **TFTP**.
- Step 6** Press the **Tab** key until you reach the local IP Address field.
- Step 7** Enter the local IP address for the switch, and press the **Tab** key.
- Step 8** Enter the subnet mask for the IP address, and press the **Tab** key.
- Step 9** Enter the IP address of the default gateway, and press the **Tab** key.
- Step 10** Enter the IP address of the TFTP server, and press the **Tab** key.
- Step 11** Enter the image name (kickstart), and press the **Tab** key. This path should be relative to the TFTP server root directory.

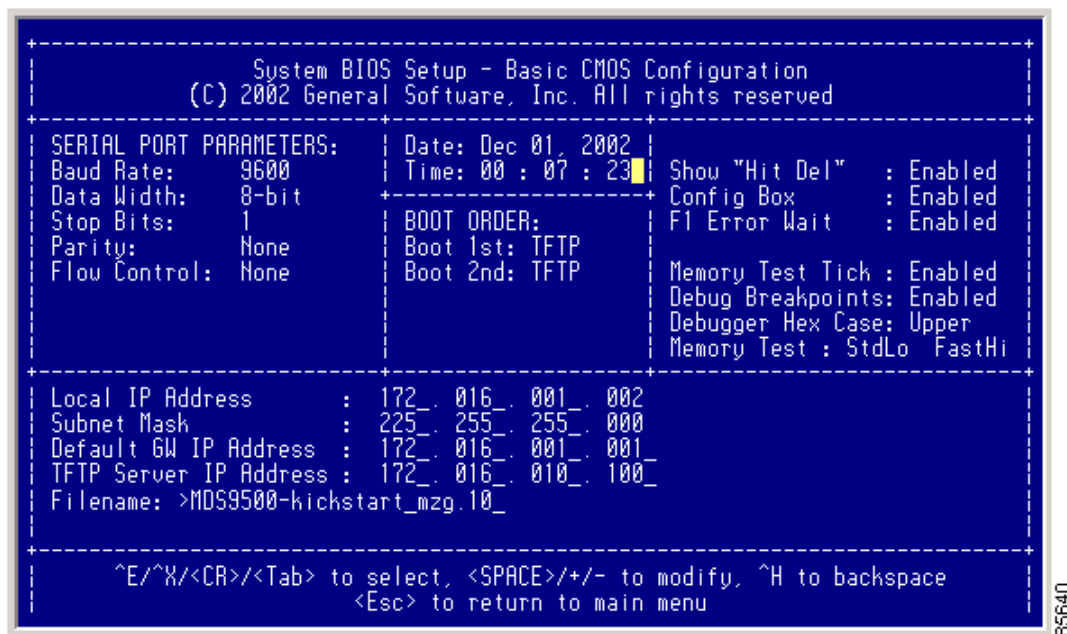


**Caution**

The file name must be entered exactly as it is displayed on your TFTP server. For example, if you have a file name **MDS9500-kickstart\_mzg.10**, then enter this name using the exact upper case characters and file extensions as shown on your TFTP server.

You see the configured changes (see [Figure 6-6](#)).

**Figure 6-6 BIOS Setup Configuration (CMOS) Changes**



**Step 12** Press the **Esc** key to return to the main menu.

**Step 13** Choose **Write to CMOS and Exit** from the main screen to save your changes.



**Note** These changes are saved in the CMOS.



**Caution** The switch must have IP connectivity to reboot using the newly configured values.

You are placed at the following prompt:

```
switch(boot)#
```

**Step 14** Enter the **init system** command at the `switch(boot)` prompt, and press **Enter**.

```
switch(boot)# init system
```

The `switch(boot)#` prompt indicates that you have a usable kickstart image.



**Note** The **init system** command also installs a new loader from the existing (running) kickstart image.

**Step 15** Follow the procedure specified in the [“Recovery from the switch\(boot\)# Prompt”](#) section on page 6-32.



## Recovery from the loader> Prompt

To recover a corrupted kickstart image (system error state) for a switch with a single supervisor module, follow these steps:

**Step 1** Press the **Esc** key to interrupt the boot loader setup after the BIOS memory test.



**Note** Press **Esc** immediately after you see the following message:

```
00000589K Low Memory Passed
00000000K Ext Memory Passed
Hit ^C if you want to run SETUP....
Wait.....
```

If you wait too long, you will skip the boot loader phase and enter the kickstart phase.

You see the loader> prompt.



**Caution**

The loader> prompt is different from the regular switch# or switch(boot)# prompt. The CLI command completion feature does not work at this prompt and may result in undesired errors. You must type the command exactly as you want the command to appear.



**Tip**

Use the **help** command at the loader> prompt to view a list of commands available at this prompt or to obtain more information about a specific command in that list.

**Step 2** Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
loader> ip address 172.16.1.2 255.255.255.0
Found Intel EtherExpressPro100 82559ER at 0xe800, ROM address 0xc000
Probing...[Intel EtherExpressPro100 82559ER]Ethernet addr: 00:05:30:00:52:27
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 0.0.0.0
```

**Step 3** Enter the IP address of the default gateway, and press **Enter**.

```
loader> ip default-gateway 172.16.1.1
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 172.16.1.1
```

**Step 4** Boot the kickstart image file from the required server, and press **Enter**.

```
loader> boot tftp://172.16.1.2/kickstart-latest
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 172.16.10.100
Gateway: 172.16.1.1
Booting: /kick-282 console=ttyS0,9600n8nn quiet loader_ver= "1.0(2)"....
.....Image verification OK
Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
```

```
INIT: Sending processes the TERM signal
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
INIT: Sending processes the TERM signal
switch(boot)#
```

The `switch(boot)#` prompt indicates that you have a usable Kickstart image.

- Step 5** Copy the system and kickstart images again.

```
switch(boot)# copy scp://user@172.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server.....

switch(boot)# copy scp://user@172.16.10.100/kickstart-img bootflash:kickstart-img
Trying to connect to tftp server.....
```

- Step 6** Follow the procedure specified in the “[Recovery from the switch\(boot\)# Prompt](#)” section on page 6-32.
- 

## Recovery from the switch(boot)# Prompt

To recover a system image using the kickstart image for a switch with a single supervisor module, follow these steps:

- Step 1** Follow this step if you issued a **init system** command. Otherwise, skip to [Step 2](#).

- a.** Change to configuration mode and configure the IP address of the switch’s `mgmt0` interface.

```
switch(boot)# config t
switch(boot) (config)# interface mgmt0
```

- b.** Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
switch(boot) (config-mgmt0)# ip address 172.16.1.2 255.255.255.0
```

- Step 2** Issue the **no shut** command to enable the interface on the switch, and press **Enter**.

```
switch(boot) (config-mgmt0)# no shut
```

- Step 3** Follow this step if you issued a **init system** command. Otherwise, skip to [Step 4](#).

- a.** Enter the IP address of the default gateway, and press **Enter**.

```
switch(boot) (config-mgmt0)# ip default-gateway 172.16.1.1
```

- Step 4** Exit to EXEC mode.

```
switch(boot) (config-mgmt0)# end
```

- Step 5** Copy the system image from the required TFTP server, and press **Enter**.

```
switch(boot)# copy scp://user@172.16.10.100/system-img bootflash:system-img
```

- Step 6** Copy the kickstart image from the required TFTP server, and press the **Enter** key.

```
switch(boot)# copy scp://user@172.16.10.100/kickstart-img bootflash:kickstart-img
```

**Step 7** Verify that the system and kickstart image files are copied to your bootflash: directory.

```
switch boot)# dir bootflash:
40295206 Aug 05 15:23:51 1980 ilc1.bin
12456448 Jul 30 23:05:28 1980 kickstart-image1
12288 Jun 23 14:58:44 1980 lost+found/
27602159 Jul 30 23:05:16 1980 system-image1
12447232 Aug 05 15:08:30 1980 kickstart-image2
28364853 Aug 05 15:11:57 1980 system-image2

Usage for bootflash://sup-local
 135404544 bytes used
 49155072 bytes free
184559616 bytes total
```

**Step 8** Load the system image from the bootflash: directory:

```
switch boot)# load bootflash:system-image
Uncompressing system image: bootflash:/system-image
CC

Would you like to enter the initial configuration mode? (yes/no): yes
```

See the [“Initial Setup Routine” section on page 4-2](#).



**Note** If you enter **no** at this point, you will return to the `switch#` login prompt, and you must manually configure the switch.

## Recovery for Switches with Dual Supervisor Modules

If one supervisor module is functioning and the other is not, boot the functioning supervisor module. Then use the booted supervisor module to bring up the supervisor module that is stuck. Issue the **reload module slot force-dnld** command (after you log into the switch) where *slot* is the slot number of the stuck supervisor module.

If both supervisor modules are not functioning, treat it like a single supervisor module recovery. First recover the image on one supervisor module and then follow the single supervisor recovery process.

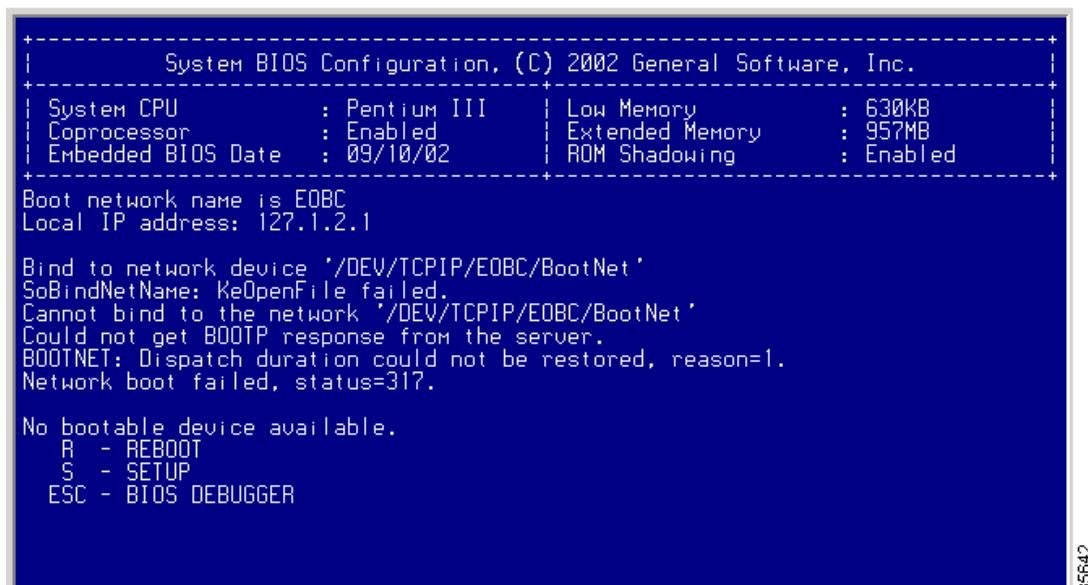


**Note** If you do not issue the **reload module** command when a boot failure has occurred, the active supervisor module automatically reloads the standby supervisor module within 3 to 6 minutes after the failure (see the [“Standby Supervisor Boot Alert” section on page 6-23](#)).

## Recognizing Error States

If you see the error messages displayed in [Figure 6-7](#) or [Figure 6-8](#), follow the procedure specified in the “Recovery Using BIOS Setup” section on page 6-27.

**Figure 6-7 Error State if Powered On and Ctrl-C is Entered**



```

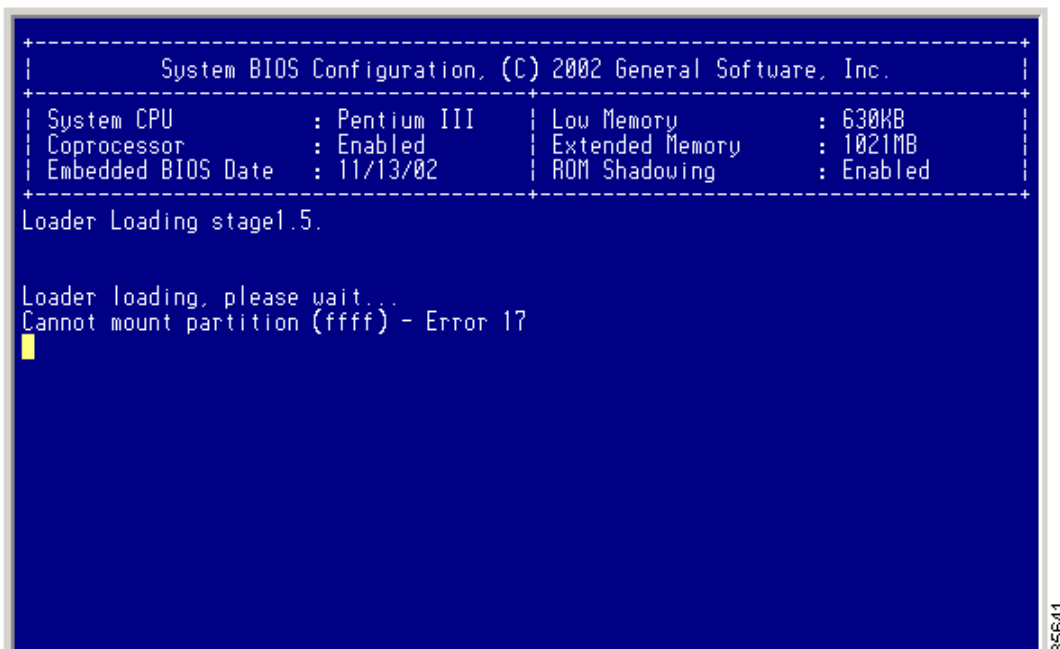
+-----+
| System BIOS Configuration, (C) 2002 General Software, Inc. |
+-----+
System CPU : Pentium III	Low Memory : 630KB
Coprocessor : Enabled	Extended Memory : 957MB
Embedded BIOS Date : 09/10/02	ROM Shadowing : Enabled
+-----+
Boot network name is EOBC
Local IP address: 127.1.2.1

Bind to network device '/DEV/TCPIP/EOBC/BootNet'
SoBindNetName: KeOpenFile failed.
Cannot bind to the network '/DEV/TCPIP/EOBC/BootNet'
Could not get BOOTP response from the server.
BOOTNET: Dispatch duration could not be restored, reason=1.
Network boot failed, status=317.

No bootable device available.
R - REBOOT
S - SETUP
ESC - BIOS DEBUGGER

```

**Figure 6-8 Error State if Powered On and Esc is Pressed**



```

+-----+
| System BIOS Configuration, (C) 2002 General Software, Inc. |
+-----+
System CPU : Pentium III	Low Memory : 630KB
Coprocessor : Enabled	Extended Memory : 1021MB
Embedded BIOS Date : 11/13/02	ROM Shadowing : Enabled
+-----+
Loader Loading stage1.5.

Loader loading, please wait...
Cannot mount partition (ffff) - Error 17

```

# Default Factory Settings

Table 6-3 lists the default settings for all Cisco MDS 9000 Family switches.

**Table 6-3 Default Factory Settings**

| Parameters      | Default               |
|-----------------|-----------------------|
| Kickstart image | No image is specified |
| System image    | No image is specified |





# Managing Modules

This chapter describes how to manage switching modules (also known as line cards) and provides information on monitoring module states. This chapter includes the following sections:

- [About Modules, page 7-1](#)
- [Verifying the Status of a Module, page 7-2](#)
- [Viewing the State of a Module, page 7-3](#)
- [Connecting to a Module, page 7-4](#)
- [Reloading Modules, page 7-5](#)
- [Preserving Module Configuration, page 7-6](#)
- [Purging Module Configuration, page 7-7](#)
- [Powering Off Switching Modules, page 7-7](#)
- [Identifying Module LEDs, page 7-8](#)
- [Configuring EPLDs, page 7-10](#)
- [Default Supervisor Module Settings, page 7-12](#)

## About Modules

[Table 7-1](#) describes the supervisor module options for switches in the Cisco MDS 9000 Family.

**Table 7-1 Supervisor Module Options**

| Product        | No. of Supervisor Modules                    | Supervisor Module Slot | Switching Module Features                                              |
|----------------|----------------------------------------------|------------------------|------------------------------------------------------------------------|
| Cisco MDS 9216 | One module (includes 16 Fibre Channel ports) | 1                      | 2-slot chassis allows one optional switching module in the other slot. |
| Cisco MDS 9509 | Two modules                                  | 5 and 6                | 9-slot chassis allows any switching module in the other seven slots.   |
| Cisco MDS 9506 | Two modules                                  | 5 and 6                | 6-slot chassis allows any switching module in the other four slots.    |

## Supervisor Modules

Supervisor modules are automatically powered up and started with the switch.

Cisco MDS 9200 Series switches have one supervisor module that includes an integrated 16-port switching module.

Cisco MDS 9500 Series switches have two supervisor modules—one in slot 5 (sup-1) and one in slot 6 (sup-2). When the switch powers up and both supervisor modules come up together, the module that enters the active mode is dependent on which of the two modules comes up first. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

## Switching Modules

Cisco MDS 9000 Family switches support any switching module in any non-supervisor slot. The switching module obtains its image from the supervisor module.

The interfaces in each module are ready to be configured when the **ok** status is displayed in the **show module** command output (see the [“Configuring Fibre Channel Interfaces”](#) section on page 10-2).

## Verifying the Status of a Module

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
```

| Mod | Ports | Module-Type                | Model           | Status     |
|-----|-------|----------------------------|-----------------|------------|
| 2   | 8     | IP Storage Services Module | DS-X9308-SMIP   | ok         |
| 4   | 0     | Caching Services Module    |                 | powered-dn |
| 5   | 0     | Supervisor/Fabric-1        | DS-X9530-SF1-K9 | active *   |
| 6   | 0     | Supervisor/Fabric-1        | DS-X9530-SF1-K9 | ha-standby |
| 8   | 0     | Caching Services Module    | DS-X9560-SMAP   | ok         |
| 9   | 32    | 1/2 Gbps FC Module         | DS-X9032        | ok         |

| Mod | Sw          | Hw    | World-Wide-Name(s) (WWN)                           |
|-----|-------------|-------|----------------------------------------------------|
| 2   | 1.3(0.106a) | 0.206 | 20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00 |
| 5   | 1.3(0.106a) | 0.602 | --                                                 |
| 6   | 1.3(0.106a) | 0.602 | -- <----- New running version in module 6          |
| 8   | 1.3(0.106a) | 0.702 | --                                                 |
| 9   | 1.3(0.106a) | 0.3   | 22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00 |

| Mod | MAC-Address(es)                        | Serial-Num  |
|-----|----------------------------------------|-------------|
| 2   | 00-05-30-00-9d-d2 to 00-05-30-00-9d-de | JAB064605a2 |
| 5   | 00-05-30-00-64-be to 00-05-30-00-64-c2 |             |
| 6   | 00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd | JAB06350B1R |
| 8   | 00-05-30-01-37-7a to 00-05-30-01-37-fe | JAB072705ja |
| 9   | 00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 | JAB06280ae9 |

\* this terminal session



The Status column in the output should display an ok status for switching modules and an active or standby (or HA-standby) status for supervisor modules. If the status is either ok or active, you can continue with your configuration.

**Note**

A standby supervisor module reflects the HA-standby status if the HA switchover mechanism is enabled (see the [“HA Switchover” section on page 5-3](#)). If the warm switchover mechanism is enabled, the standby supervisor module reflects the standby status.

The states through which a switching module progresses is discussed in the [“Viewing the State of a Module” section on page 7-3](#).

## Viewing the State of a Module

If your chassis has more than one switching module (also known as line card), you will see the progress check if you issue the **show module** command several times and view the status column each time.

The switching module goes through a testing and an initializing stage before displaying an ok status. [Table 7-2](#) describes the possible states in which a module can exist.

**Table 7-2** Module States

| show module Output | Description                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| powered up         | The hardware has electrical power. When the hardware is powered up, the software begins booting.                                                                                                                                                          |
| testing            | The module has established connection with the supervisor and the switching module is performing bootup diagnostics.                                                                                                                                      |
| initializing       | The diagnostics have completed successfully and the configuration is being downloaded.                                                                                                                                                                    |
| failure            | The switch detects a switching module failure upon initialization and automatically attempts to power-cycle the module three times. After the third attempt it continues to display a failed state.                                                       |
| ok                 | The switch is ready to be configured.                                                                                                                                                                                                                     |
| power-denied       | The switch detects insufficient power for a switching module to power up.<br>In this case, issue a <b>show environment power</b> command to determine power consumption issues (see <a href="#">Chapter 31, “Monitoring System Processes and Logs”</a> ). |
| active             | This module is the active supervisor module and the switch is ready to be configured.                                                                                                                                                                     |
| HA-standby         | This module is the standby supervisor module and that the HA switchover mechanism is enabled (see the <a href="#">“HA Switchover” section on page 5-3</a> ).                                                                                              |
| standby            | This module is the standby supervisor module and the warm switchover mechanism is enabled (see the <a href="#">“HA Switchover” section on page 5-3</a> ).                                                                                                 |

## Connecting to a Module

At any time, you can connect to any module using the **attach module** command. Once you are at the module prompt, you can obtain further details about the module using module-specific commands in EXEC mode.

To attach to a module, follow these steps:

|        | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Purpose                                                                                                                                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>attach module 6</b><br>switch(standby) #                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Provides direct access to the specified module (in this example, the standby supervisor module is in slot 6).                                                                                                                                                                                    |
| Step 2 | switch(standby) # <b>dir bootflash:</b><br><br><pre> root      14502912   Jan 13 12:23:52 1980 kickstart_image1 admin     14424576   Jan 14 06:47:29 1980 kickstart_image2 admin     14469632   Jan 14 01:29:16 1980 kickstart_image3 root      14490112   Jan 08 07:25:50 1980 kickstart_image4 root       12288     Jan 16 15:49:24 1980 lost+found/ admin     14466048   Jan 14 02:40:16 1980 kickstart_image5 admin     24206675   Jan 14 02:57:03 1980 m9500-sflek.bin root      19084510   Jan 13 12:23:28 1980 system_image1 admin     19066505   Jan 14 06:45:16 1980 system_image2 admin     18960567   Jan 14 01:25:21 1980 system_image5            Usage for bootflash: filesystem                     158516224 bytes total used                       102400 bytes free                     167255040 bytes available </pre> | Provides the available space information for the standby supervisor module.<br><br><b>Note</b> Type <b>exit</b> to exit the module-specific prompt.<br><br><b>Tip</b> If you are not accessing the switch from a console terminal, this is the only way to access the standby supervisor module. |

You can also use the **attach module** command as follows:

- To view the standby supervisor module information, but you can not configure the standby supervisor module using this command.
- On the switching module portion of the Cisco MDS 9216 supervisor module which resides in slot 1.

# Reloading Modules

You can reload the entire switch, reset specific modules in the switch, or reload the image on specific module in the switch.

## Reloading the Switch

To reload the switch, issue the **reload** command without any options. When you issue this command, you reboot the switch (see [Chapter 6, “Software Images”](#)).

**Note**

If you need to issue the **reload** command, be sure to save the running configuration using the **copy running-config startup-config** command.

## Power Cycling Modules

To power cycle any module, follow these steps:

- 
- Step 1** Identify the module that needs to be reset.
- Step 2** Issue the **reload module** command to reset the identified module. This command merely power cycles the selected module.

```
switch# reload module number
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch# reload module 2
```

---

## Reloading Switching Modules

Switching modules automatically download their images from the supervisor module, and do not need a force download. This procedure is provided for reference should a need arise.

To replace the image on a switching module, follow these steps:

- 
- Step 1** Identify the switching module that requires the new image.
- Step 2** Issue the **reload module number force-dnld** command to update the image on the switching module.

```
switch# reload module number force-dnld
```

Where *number* indicates the slot in which the identified module resides. For example, if the identified module resides in slot 9:

```
switch# reload module 9 force-dnld...
Jan 1 00:00:46 switch %LC-2-MSG:SLOT9 LOG_LC-2-IMG_DNLD_COMPLETE: COMPLETED
downloading of linecard image. Download successful...
```

---

# Preserving Module Configuration

To save the configuration, enter the **copy running-config startup-config** command from the EXEC mode prompt to save the new configuration into nonvolatile storage. Once this command is issued, the running and the startup copies of the configuration are identical.

Table 7-3 displays various scenarios when module configurations are persevered or lost.

**Table 7-3 Switching Module Configuration Status**

| Scenario                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Consequence                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A particular switching module is removed and the <b>copy running-config startup-config</b> command is issued again.                                                                                                                                                                                                                                                                                                                                                                 | The configured module information is lost.                                                                                                                                                                                                                                                                                                                                                                |
| A particular switching module is removed and the same switching module is replaced before the <b>copy running-config startup-config</b> command is issued again.                                                                                                                                                                                                                                                                                                                    | The configured module information is preserved.                                                                                                                                                                                                                                                                                                                                                           |
| A particular switching module is removed and replaced with the same type switching module, and a <b>reload module number</b> command is issued.                                                                                                                                                                                                                                                                                                                                     | The configured module information is preserved.                                                                                                                                                                                                                                                                                                                                                           |
| A particular switching module is removed and replaced with a different type of switching module. For example, a 16-port switching module is replaced with a 32-port switching module.                                                                                                                                                                                                                                                                                               | The configured module information is lost from the running configuration. The default configuration is applied.<br><br>The configured module information remains in startup configuration until a <b>copy running-config startup-config</b> command is issued again.                                                                                                                                      |
| <p>Sample scenario:</p> <ol style="list-style-type: none"> <li>1. The switch currently has a 16-port switching module and the startup and running configuration files are the same.</li> <li>2. You replace the 16-port switching module in the switch with a 32-port switching module.</li> <li>3. Next, you remove the 32-port switching module and replace it with the same 16-port switching module referred to in Step 1.</li> <li>4. You <b>reload</b> the switch.</li> </ol> | <p>Sample response:</p> <ol style="list-style-type: none"> <li>1. The switch uses the 16-port switching module and the present configuration is saved in nonvolatile storage.</li> <li>2. The factory default configuration is applied.</li> <li>3. The factory default configuration is applied.</li> <li>4. The configuration saved in nonvolatile storage referred to in Step 1 is applied.</li> </ol> |

## Purging Module Configuration

To delete the configuration in a specific module, issue the **purge module *slot* running-config** command from EXEC mode. Once this command is issued, the running configuration is cleared for the specified slot. This command will not work on supervisor modules or on any slot which currently has a module. This command only works on an empty slot (where the specified module once resided).

The **purge module** command clears the configuration for any module that previously existed in a slot and has since been removed. While the module was in that slot, some parts of the configuration may have been stored in the running configuration and cannot be reused (for example, IP addresses), unless it is cleared from the running configuration.

For example, if you create an IP storage configuration with a an IPS module in slot 3 in Switch A. This module uses the 10.1.5.500 IP address. You decide to remove this IPS module and move it to Switch B, and you no longer need the 10.1.5.500 IP address. If you try to configure this unused 10.1.5.500 IP address, you will receive an error message that prevents you from proceeding with the configuration. In this case, you need to issue the **purge module 3 running-config** command to clear the old configuration in Switch A and then proceed with using this IP address.

## Powering Off Switching Modules

By default, all switching modules are configured to be in the power up state.

To power off a module, follow these steps:

|        | Command                                                        | Purpose                                                             |
|--------|----------------------------------------------------------------|---------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                        | Enters configuration mode.                                          |
| Step 2 | switch(config)# <b>poweroff module 1</b><br>switch(config)#    | Powers off the specified module (switching module 1) in the switch. |
|        | switch(config)# <b>no poweroff module 1</b><br>switch(config)# | Powers up the specified module (switching module 1) in the switch.  |

# Identifying Module LEDs

Table 7-4 to Table 7-7 describe the LED location, type, and status for supervisor and switching modules used in Cisco MDS 9000 Family switches.

**Table 7-4** *Module LEDs on a Cisco MDS 9200 Series Switch*

| Module                    | LED Type | Status | Description                                                                                                                                                                                                                                                                                                                       |
|---------------------------|----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fixed switching module    | Status   | Green  | <ul style="list-style-type: none"> <li>All chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) are reporting OK.</li> <li>Sufficient power is available for all modules</li> </ul>                                                                                                |
|                           |          | Orange | <ul style="list-style-type: none"> <li>Any one of the chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) failed.</li> <li>Sufficient power is not available for all modules.</li> <li>Incompatible power supplies are installed.</li> <li>The redundant clock failed.</li> </ul> |
|                           |          | Red    | <ul style="list-style-type: none"> <li>The diagnostic test failed.</li> <li>The module is not operational because a fault occurred during the initialization sequence.</li> <li>A temperature condition occurred. (A major threshold was exceeded during environmental monitoring.)</li> </ul>                                    |
| Optional switching module | System   | Green  | All diagnostics pass. The module is operational (normal initialization sequence).                                                                                                                                                                                                                                                 |
|                           |          | Orange | <ul style="list-style-type: none"> <li>The module is booting or running diagnostics (normal initialization sequence).</li> <li>An over temperature condition occurred. (A minor threshold was exceeded during environmental monitoring.)</li> </ul>                                                                               |
|                           |          | Red    | <ul style="list-style-type: none"> <li>The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence.</li> <li>An over temperature condition occurred. (A major threshold was exceeded during environmental monitoring.)</li> </ul>                                       |

**Table 7-5 Supervisor Module LEDs on a Cisco MDS 9500 Series Switch**

| LED                   | Status | Description                                                                                                                                                                                                                                                                                          |
|-----------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status                | Green  | All diagnostics pass. The module is online.                                                                                                                                                                                                                                                          |
|                       | Orange | <ul style="list-style-type: none"> <li>The module is booting or running diagnostics (normal initialization sequence).</li> <li>The module is not online.</li> <li>An over temperature condition has occurred. (A minor threshold has been exceeded during environmental monitoring.)</li> </ul>      |
|                       | Red    | <ul style="list-style-type: none"> <li>The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence.</li> <li>An over temperature condition has occurred. (A major threshold has been exceeded during environmental monitoring.)</li> </ul> |
| System <sup>1</sup>   | Green  | All chassis environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) are reporting OK.                                                                                                                                                                            |
|                       | Orange | <ul style="list-style-type: none"> <li>Any one of the environmental monitors (power supply, fan, temperature sensor, clock, and chassis SEE PROM) has failed.</li> <li>Incompatible power supplies are installed.</li> <li>The redundant clock has failed.</li> </ul>                                |
|                       | Red    | The temperature of the supervisor module major threshold has been exceeded.                                                                                                                                                                                                                          |
| Active                | Green  | The supervisor module is operational and active.                                                                                                                                                                                                                                                     |
|                       | Orange | The supervisor module is in standby mode.                                                                                                                                                                                                                                                            |
| Pwr Mgmt <sup>1</sup> | Green  | Sufficient power is available for all modules.                                                                                                                                                                                                                                                       |
|                       | Orange | Sufficient power is not available for all modules.                                                                                                                                                                                                                                                   |

1. The System and Pwr Mgmt LEDs on a redundant supervisor module are synchronized to the active supervisor module.

**Table 7-6 Ethernet Interface LEDs on a Cisco MDS 9200 Series Switch**

| Module            | LED Type | Status         | Description                               |
|-------------------|----------|----------------|-------------------------------------------|
| Ethernet (mgmt 0) | Activity | Flashing green | Traffic is passing through the interface. |
|                   | Link     | Solid green    | The link is functioning.                  |
|                   |          | Off            | The link is down.                         |

**Table 7-7 Switching Module LEDs**

| LED Type | Status          | Description                                                                                                                                                                                                                                                                                 |
|----------|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status   | Green           | All diagnostics pass. The module is operational (normal initialization sequence).                                                                                                                                                                                                           |
|          | Orange          | <ul style="list-style-type: none"> <li>The module is booting or running diagnostics (normal initialization sequence).</li> <li>An over temperature condition occurred. (A minor threshold was exceeded during environmental monitoring.)</li> </ul>                                         |
|          | Red             | <ul style="list-style-type: none"> <li>The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence.</li> <li>An over temperature condition occurred. (A major threshold was exceeded during environmental monitoring.)</li> </ul> |
| Speed    | On              | 2 Gbps mode.                                                                                                                                                                                                                                                                                |
|          | Off             | 1 Gbps mode.                                                                                                                                                                                                                                                                                |
| Link     | Solid green     | Link is up.                                                                                                                                                                                                                                                                                 |
|          | Flashing green  | Link is up (beacon used to identify port). See the <a href="#">“Identifying the Beacon LEDs”</a> section on page 10-13.                                                                                                                                                                     |
|          | Solid yellow    | Disabled by software.                                                                                                                                                                                                                                                                       |
|          | Flashing yellow | Fault is detected.                                                                                                                                                                                                                                                                          |
|          | Off             | Link is down.                                                                                                                                                                                                                                                                               |

## Configuring EPLDs

Switches and directors in the Cisco MDS 9000 Family contain several electrically programmable logical devices (EPLDs) that provide hardware functionalities in all modules. Starting with Cisco MDS SAN-OS Release 1.2, EPLD image upgrades are periodically provided to include enhanced hardware functionality or to resolve known issues.



### Tip

Refer to the Cisco MDS SAN-OS Release Notes to verify if the EPLD has changed for the SAN-OS image version being used.

EPLDs can be upgraded or downgraded using CLI commands. When EPLDs are being upgraded or downgraded, the following guidelines and observations apply:

- You can individually update each module that is online. The EPLD update is only disruptive to the module being upgraded.
- If you interrupt an upgrade, the module must be upgraded again.
- The upgrade or downgrade can only be executed from the active supervisor module. While the active supervisor module cannot be updated, you can update the other modules individually.
- In Cisco MDS 9100 Series Fabric switches, be sure to specify 1 as the module number.
- Cisco MDS 9216 Switches do not support EPLD upgrades.



**Caution**

Do not insert or remove any modules while an EPLD upgrade or downgrade is in progress.

To upgrade EPLD images for a module, issue the **install module *number* epld *url*** command on the active supervisor module:

```
switch# install module 2 epld scp://user@10.6.16.22/m9000-epld-version.img

The authenticity of host '10.6.16.22' can't be established.
RSA1 key fingerprint is 55:2e:1f:0b:18:76:24:02:c2:3b:62:dc:9b:6b:7f:b7.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.6.16.22' (RSA1) to the list of known hosts.
user@10.6.16.22's password:
epld.img 100% |*****| 1269 KB 00:00

Module Number 2
EPLD Curr Ver New Ver

Power Manager 0x06
XBUS IO 0x07 0x08
UD Flow Control 0x05
PCI ASIC I/F 0x05 0x05

Module 2 will be powered down now!!
Do you want to continue (y/n) ? y
\ <-----progress twirl
Module 2 EPLD upgrade is successful
```

If you forcefully upgrade a module that is not online, all EPLDs are forcefully upgraded. If the module is not present in the switch, an error is returned. If the module is present, the command process continues. To update a module that is not online but is present in the chassis, use the same command. The switch software prompts you to continue after reporting the module state. When you confirm your intention to continue, the upgrade continues

```
switch# install module 2 epld scp://user@10.6.16.22/m9000-epld-version.img

Module 2 is not online, Do you want to continue (y/n) ? y
cchetty@171.69.16.22's password:
epld.img 100% |*****| 1269 KB 00:00
\ <-----progress twirl
Module 2 EPLD upgrade is successful
```

**Note**

Switches in the Cisco MDS 9100 Series do not support a forced EPLD upgrade. When you upgrade the EPLD module on these switches, you receive the following message:

```
Data traffic on the switch will stop now!!
Do you want to continue (y/n) ?
```

## Displaying EPLD Versions

To view all current EPLD versions on a specified module, use the **show version module *number* epld** command (see [Example 7-1](#)).

### **Example 7-1** Displays Current EPLD Versions for a Specified Module

```
switch# show version module 2 epld
Module Number 2
```

| EPLD Device     | Version |
|-----------------|---------|
| Power Manager   | 0x06    |
| XBUS IO         | 0x07    |
| UD Flow Control | 0x05    |
| PCI ASIC I/F    | 0x05    |

To view the available EPLD versions, use the **show version epld url** command (see [Example 7-2](#)).

### Example 7-2 Displays Available EPLD Versions

```
switch# show version epld scp://user@10.6.16.22/m9000-epld-version.img
user@10.6.16.22's password:
```

| Module Name                   | EPLD Device             | Version |
|-------------------------------|-------------------------|---------|
| MDS 9500 Supervisor 1         | XBUS 1 IO               | 0x09    |
|                               | XBUS 2 IO               | 0x0c    |
|                               | UD Flow Control         | 0x05    |
|                               | PCI ASIC I/F            | 0x04    |
| 1/2 Gbps FC Module (16 port)  | XBUS IO                 | 0x08    |
|                               | PCI ASIC I/F            | 0x05    |
| 1/2 Gbps FC Module (32 port)  | XBUS IO                 | 0x07    |
|                               | PCI ASIC I/F            | 0x05    |
| Advanced Services Module      | XBUS IO                 | 0x07    |
|                               | UD Flow Control         | 0x05    |
|                               | PCI Bridge              | 0x04    |
| IP Storage Services Module    | XBUS IO                 | 0x02    |
|                               | UD Flow Control         | 0x05    |
|                               | PCI ASIC I/F            | 0x05    |
|                               | Services Module I/F0x12 |         |
|                               | IPS DB I/F              | 0x08    |
| MDS 9100 Series Fabric Switch | XBUS IO                 | 0x03    |

## Default Supervisor Module Settings

[Table 7-8](#) lists the default settings for the supervisor module.

**Table 7-8 Default Supervisor Module Settings**

| Parameters                 | Default                                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Administrative connection  | Serial connection.                                                                                                                                |
| Global switch information  | <ul style="list-style-type: none"> <li>No value for system name.</li> <li>No value for system contact.</li> <li>No value for location.</li> </ul> |
| System clock               | No value for system clock time.                                                                                                                   |
| In-band (VSAN 1) interface | IP address, subnet mask, and broadcast address assigned to the VSAN is set to 0.0.0.0.                                                            |



## Managing System Hardware

---

This chapter provides details on monitoring the health of the switch. It includes the following sections:

- [Displaying Switch Hardware Inventory, page 8-2](#)
- [Displaying Power Usage Information, page 8-5](#)
- [Displaying Power Usage Information, page 8-5](#)
- [Configuring Power Supplies, page 8-6](#)
- [Displaying Module Temperature, page 8-9](#)
- [Monitoring Fan Modules, page 8-10](#)
- [Monitoring Clock Modules, page 8-10](#)
- [Displaying Environment Information, page 8-11](#)

# Displaying Switch Hardware Inventory

Use the **show hardware** command to display switch hardware inventory details. See [Example 8-1](#).



## Note

To display and configure modules, see [Chapter 7, “Managing Modules.”](#)

### Example 8-1 Displays the Hardware Information

```
switch# show hardware
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license.

Software
 BIOS: version 1.0.3
 loader: version error [last 1.0(1)]
 kickstart: version 1.1(1) [build 1.1(0.94)] [gdb]
 system: version 1.1(1) [build 1.1(0.94)] [gdb]

 BIOS compile time: 11/18/02
 kickstart image file is: bootflash:/bootimage
 kickstart compile time: 2/12/2003 11:00:00
 system image file is: isanimage
 system compile time: 2/12/2003 12:00:00

Hardware
 RAM 1027628 kB

 bootflash: 1000944 blocks (block size 512b)
 slot0: 0 blocks (block size 512b)

 172.22.90.171 uptime is 0 days 3 hours 0 minute(s) 36 second(s)

 Last reset at 669882 usecs after Thu Feb 13 07:20:41 2003
 Reason: Reset Requested by CLI command reload
 System version: 1.0(1)

This supervisor carries Pentium processor with 1027628 kB of memory
Intel(R) Pentium(R) III CPU at family with 512 KB L2 Cache
Rev: Family 6, Model 11 stepping 1

512K bytes of non-volatile memory.
1000944 blocks of internal bootflash (block size 512b)

Chassis has 9 slots for Modules

Module in slot 1 is empty

Module in slot 2 is empty

Module in slot 3 is ok
 Module type is "1/2 Gbps FC Module"
 1 submodules are present
 RAM size is 0 (kb)
 Model number is DS-X9016
```

```
H/W version is 0.0
Part Number is 73-8127-03
Part Revision is
Manufacture Date is Year 0 Week 0
Serial number is
CLEI code is

Module in slot 4 is empty

Module in slot 5 is ok
 Module type is "Supervisor/Fabric-1"
 No submodules are present
 Model number is DS-X9530-SF1-K9
 H/W version is 0.0
 Part Number is 73-7523-06
 Part Revision is
 Manufacture Date is Year 0 Week 0
 Serial number is
 CLEI code is

Module in slot 6 is empty

Module in slot 7 is empty

Module in slot 8 is ok
 Module type is "IP Storage Module"
 3 submodules are present
 RAM size is 0 (kb)
 Model number is DS-X9308-SMIP
 H/W version is 0.2
 Part Number is 73-8083-02
 Part Revision is 2
 Manufacture Date is Year 7 Week 2
 Serial number is JAB0702065h
 CLEI code is 0

Module in slot 9 is ok
 Module type is "1/2 Gbps FC Module"
 1 submodules are present
 RAM size is 0 (kb)
 Model number is DS-X9016
 H/W version is 0.0
 Part Number is 73-8127-03
 Part Revision is
 Manufacture Date is Year 0 Week 0
 Serial number is
 CLEI code is

Chassis has 2 Slots for Power Supplies

PS in slot A is ok
 Power supply type is "1153.32W 110v AC"
 Model number is DS-CAC-2500W
 H/W version is 1.0
 Part Number is 341-0061-01
 Part Revision is A0
 Manufacture Date is Year 6 Week 16
 Serial number is ART061600VA
 CLEI code is

PS in slot B is ok
 Power supply type is "1153.32W 110v AC"
```

```

Model number is WS-CAC-2500W
H/W version is 1.0
Part Number is 34-1535-01
Part Revision is A0
Manufacture Date is Year 6 Week 20
Serial number is ART0620005N
CLEI code is

```

```

Chassis has one slot for Fan Module

```

```

Fan module is ok
Model number is WS-9SLOT-FAN
H/W version is 0.0
Part Number is 800-22342-01
Part Revision is
Manufacture Date is Year 0 Week 0
Serial number is
CLEI code is

```

## Displaying the Switch Serial Number

The serial number of your Cisco MDS 9000 Family switch can be obtained by looking at the serial number label on the back of the switch (next to the power supply), or by executing the operating system **show sprom backplane 1** command.

```

switch# show sprom backplane 1
DISPLAY backplane sprom contents:
Common block:
Block Signature : 0xabab
Block Version : 2
Block Length : 156
Block Checksum : 0x106f
EEPROM Size : 512
Block Count : 3
FRU Major Type : 0x6001
FRU Minor Type : 0x0
OEM String : Cisco Systems, Inc.
Product Number : DS-C9506
Serial Number : FOX0712S007
Part Number : 73-8697-01
Part Revision : 01
Mfg Deviation : 0
H/W Version : 0.1
Mfg Bits : 0
Engineer Use : 0
snmpOID : 9.12.3.1.4.26.0.0
Power Consump : 0
RMA Code : 0-0-0-0
Chassis specific block:
...

```

# Displaying Power Usage Information

Use the **show environment power** command to display the actual power usage information for the entire switch. In response to this command, power supply capacity and consumption information is displayed for each module. See [Example 8-2](#).



## Note

In a Cisco MDS 9500 Series switch, power usage is reserved for both supervisors regardless of whether one or both supervisor modules are present.

### Example 8-2 Displays Power Management Information

```
switch# show environment power
```

```

PS Model Power Power Status
 (Watts) (Amp @42V)

1 DS-CAC-2500W 1153.32 27.46 ok
2 WS-CAC-2500W 1153.32 27.46 ok

Mod Model Power Power Power Power Status
 Requested Requested Allocated Allocated
 (Watts) (Amp @42V) (Watts) (Amp @42V)

1 DS-X9032 199.92 4.76 199.92 4.76 powered-up
4 DS-X9032 199.92 4.76 199.92 4.76 powered-up
5 DS-X9530-SF1-K9 126.00 3.00 126.00 3.00 powered-up
6 DS-X9530-SF1-K9 126.00 3.00 126.00 3.00 powered-up
9 DS-X9016 220.08 5.24 220.08 5.24 powered-up

Power Usage Summary:

Power Supply redundancy mode: redundant

Total Power Capacity 1153.32 W

Power reserved for Supervisor(s) [-] 252.00 W
Power reserved for Fan Module(s) [-] 0.00 W
Power currently used by Modules[-] 619.92 W

Total Power Available 281.40 W

```

# Configuring Power Supplies

Switches in the MDS 9000 Family have two redundant power supply slots. The power supplies can be configured in either **redundant** or **combined** mode.

- **redundant**—Uses the capacity of one power supply only. This is the default mode. In case of power supply failure, the entire switch has sufficient power available in the system.
- **combined**—Uses the combined capacity of both power supplies. In case of power supply failure, the entire switch can be shut down (depends on the power used) causing traffic disruption. This mode is seldom used, except in cases where the switch has two low power supply capacities but a higher power usage.



## Note

The chassis in the Cisco MDS 9000 Family uses 1200Watts when powered at 110 volts, and 2500Watts when powered at 220 volts.

To configure the power supply mode, follow these steps:

|        | Command                                                                   | Purpose                                               |
|--------|---------------------------------------------------------------------------|-------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                   | Enters configuration mode.                            |
| Step 2 | switch(config)# <b>power redundancy-mode combined</b><br>switch(config)#  | Configures combined power supply mode.                |
|        | switch(config)# <b>power redundancy-mode redundant</b><br>switch(config)# | Reverts to the redundant (default) power supply mode. |

## Power Supply Guidelines



## Note

Use the **show environment power** command to view the current power supply configuration.

Follow these guidelines when configuring power supplies:

1. When power supplies with different capacities are installed in the switch, the total power available differs based on the configured mode:
  - In **redundant** mode, the total power is the lesser of the two power supply capacities. For example, if you have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = not used

Current usage = 2000Watts

Current capacity = 2500Watts

Then the following three scenarios differ as specified (see [Table 8-1](#)):

- a. **Scenario 1:** If 1800Watts is added as power supply 2, then power supply 2 is shut down.  
Reason: 1800Watts is less than the usage of 2000Watts.
- b. **Scenario 2:** If 2200Watts is added as power supply 2, then the current capacity decreases to 2200Watts.  
Reason: 2200Watts is the lesser of the two power supplies.



- c. **Scenario 3:** If 3000Watts is added as power supply 2, then the current capacity value remains at 2500Watts.

Reason: 2500Watts is the lesser of the two power supplies.

**Table 8-1 Redundant Mode Power Supply Scenarios**

| Scenario | Power Supply 1 (W) <sup>1</sup> | Current Usage (W) | Insertion of Power Supply 2 (W) | New Capacity (W) | Action Taken by Switch       |
|----------|---------------------------------|-------------------|---------------------------------|------------------|------------------------------|
| 1        | 2500                            | 2000              | 1800                            | 2500             | Power supply 2 is shut down. |
| 2        | 2500                            | 2000              | 2200                            | 2200             | Capacity becomes 2200Watts.  |
| 3        | 2500                            | 2000              | 3300                            | 2500             | Capacity remains the same.   |

1. W = Watts

- In **combined** mode, the total power is twice the lesser of the two power supply capacities.

For example, if you have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = not used

Current Usage = 2000Watts

Current capacity = 2500Watts

Then, the following three scenarios differ as specified (see [Table 8-2](#)):

- Scenario 1:** If 1800Watts is added as power supply 2, then the capacity increases to 3600Watts.  
Reason: 3600Watts is twice the minimum (1800Watts).
- Scenario 2:** If 2200Watts is added as power supply 2, then the current capacity increases to 4400Watts.  
Reason: 4400Watts is twice the minimum (2200Watts).
- Scenario 3:** If 3000Watts is added as power supply 2, then the current capacity increases to 5000Watts.  
Reason: 5000Watts is twice the minimum (2500Watts).

**Table 8-2 Combined Mode Power Supply Scenarios**

| Scenario | Power Supply 1 (W) <sup>1</sup> | Current Usage (W) | Insertion of Power Supply 2 (W) | New Capacity (W) | Action Taken by Switch                                 |
|----------|---------------------------------|-------------------|---------------------------------|------------------|--------------------------------------------------------|
| 1        | 2500                            | 2000              | 1800                            | 3600             | Power is never shut down. The new capacity is changed. |
| 2        | 2500                            | 2000              | 2200                            | 4400             |                                                        |
| 3        | 2500                            | 2000              | 3300                            | 5000             |                                                        |

1. W = Watts

2. When you change the configuration from **combined** to **redundant** mode and the system detects a power supply that has a capacity lower than the current usage, the power supply is shut down. If both power supplies have a lower capacity than the current system usage, the configuration is not allowed. Various configuration scenarios are displayed and summarized in [Table 8-3](#).

a. **Scenario 1:** You have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = 1800Watts

Current Usage = 2000Watts

Current mode = **combined** mode (so current capacity is 3600Watts)

You decide to change the switch to **redundant** mode. Then power supply 2 is shut down.

Reason: 1800Watts is the lesser of the two power supplies and it is less than the system usage.

b. **Scenario 2:** You have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = 2200Watts

Current Usage = 2000Watts

Current mode = **combined** mode (so current capacity is 4400Watts).

You decide to change the switch to **redundant** mode. Then the current capacity decreases to 2200Watts.

Reason: 2200Watts is the lesser of the two power supplies.

c. **Scenario 3:** You have the following usage figures configured:

Power supply 1 = 2500Watts

Additional Power supply 2 = 1800Watts

Current Usage = 3000Watts

Current mode = **combined** mode (so current capacity is 3600Watts).

You decide to change the switch to **redundant** mode. Then the current capacity decreases to 2500Watts and the configuration is rejected.

Reason: 2500Watts is less than the system usage (3000Watts).

**Table 8-3 Combined Mode Power Supply Scenarios**

| Scenario | Power Supply 1 (W) <sup>1</sup> | Current Mode    | Current Usage (W) | Power Supply 2 (W) | New Mode         | New Capacity (W) | Action Taken by Switch                                 |
|----------|---------------------------------|-----------------|-------------------|--------------------|------------------|------------------|--------------------------------------------------------|
| 1        | 2500                            | <b>combined</b> | 2000              | 1800               | N/A              | 3600             | Existing configuration.                                |
|          | 2500                            | <b>N/A</b>      | 2000              | 1800               | <b>redundant</b> | 2500             | Power supply 2 is shut down                            |
| 2        | 2500                            | <b>combined</b> | 2000              | 2200               | N/A              | 4400             | Existing configuration.                                |
|          | 2500                            | <b>N/A</b>      | 2000              | 2200               | <b>redundant</b> | 2200             | The new capacity is changed.                           |
| 3        | 2500                            | <b>combined</b> | 3000              | 1800               | N/A              | 3600             | Existing configuration.                                |
|          | 2500                            | <b>N/A</b>      | 3000              | 1800               | <b>redundant</b> | N/A              | Rejected, so the mode reverts to <b>combined</b> mode. |

1. W = Watts

# Displaying Module Temperature

Use the **show environment temperature** command to display temperature sensors for each module (see [Example 8-3](#)).

Each module (switching and supervisor) has four sensors: 1 (outlet sensor), 2 (intake sensor), 3 (onboard sensor), and 4 (onboard sensor). Each sensor has two thresholds (in Celsius): minor and major.



## Note

A threshold value of -127 indicates that no thresholds are configured or applicable.

- minor threshold—When a minor threshold is exceeded, a minor alarm occurs and the following action is taken for all four sensors:
  - Syslog messages are displayed.
  - Call Home alerts are sent (if configured).
  - SNMP notifications are sent (if configured).
- major threshold—When a major threshold is exceeded, a major alarm occurs and the following action is taken as follows:
  - For sensors 1, 3, and 4 (outlet and onboard sensors):
 

Syslog messages are displayed.

Call Home alerts are sent (if configured).

SNMP notifications are sent (if configured).
  - For sensor 2 (intake sensor):
 

If the threshold is exceeded in a switching module, the module is shut down.

If the threshold is exceeded in a supervisor module with HA-standby or standby present, the supervisor module is shut down.

If the standby supervisor is not present, the entire switch is shut down.



## Note

Switch shut down only happens after a two-minute interval. During this interval the software monitors the temperature every five (5) seconds and continuously sends syslog messages as configured. If the required action is not taken (for example, a new fan module is inserted to decrease temperature) and if the temperature does not come down, the system is shut down at the end of two minutes.

### Example 8-3 Displays Temperature Information

```
switch# show environment temperature
```

| Module | Sensor | MajorThresh<br>(Celsius) | MinorThres<br>(Celsius) | CurTemp<br>(Celsius) | Status |
|--------|--------|--------------------------|-------------------------|----------------------|--------|
| 2      | Outlet | 75                       | 60                      | 35                   | ok     |
| 2      | Intake | 65                       | 50                      | 33                   | ok     |
| 5      | Outlet | 75                       | 60                      | 44                   | ok     |
| 5      | Intake | 65                       | 50                      | 36                   | ok     |
| 6      | Outlet | 75                       | 60                      | 42                   | ok     |
| 6      | Intake | 65                       | 50                      | 35                   | ok     |

|   |        |    |    |    |    |
|---|--------|----|----|----|----|
| 7 | Outlet | 75 | 60 | 33 | ok |
| 7 | Intake | 65 | 50 | 30 | ok |
| 9 | Outlet | 75 | 60 | 34 | ok |
| 9 | Intake | 65 | 50 | 39 | ok |

## Monitoring Fan Modules

Use the **show environment fan** command to display the fan status for each fan module. See [Example 8-4](#).

### Example 8-4 Displays Chassis Fan Information

```
switch# show environment fan

FAN Model Hw Status

Chassis WS-9SLOT-FAN 0.0 ok
PS-1 -- -- ok
PS-2 -- -- ok
```

The fan status is continuously monitored. In case of a fan module failure, the following action is taken:

- Syslog messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).



#### Caution

A fan failure could lead to temperature alarms if not corrected immediately.

When a fan module is removed, the `Fan module removed` message is displayed every 10 seconds for three minutes after the fan removal. After the three-minute interval, a system shutdown message is displayed every 5 seconds for two additional minutes. At the end of this 5-minute sequence, the system is shutdown. If the fan module is re-inserted at any point within this 5-minute period, the remaining sequence is stopped.

## Monitoring Clock Modules

Use the **show environment clock** command to display the clock status for the chassis. See [Example 8-5](#).

### Example 8-5 Displays Chassis Clock Information

```
switch# show environment clock

Clock Model Hw Status

A DS-C9500-CL 0.0 ok/active
B DS-C9500-CL 0.0 ok/standby
```

Each switch has two clock modules for redundancy: Clock A (primary) and Clock B. The redundant clock module (Clock B) takes over if the primary clock module fails. If Clock A is available at startup, the switch uses Clock A, otherwise it uses Clock B.

If Clock A fails, the switch is reset and Clock B automatically takes over. Clock modules cannot be configured. If both modules fail, the switch shuts down. The probability of a clock failure is low given that the mean time between failures (MTBF) is 3660316 hours.

## Displaying Environment Information

Use the **show environment** command to display all environment-related switch information.

### Example 8-6 Displays All Environment Information

```
switch# show environment
Clock:

Clock Model Hw Status

A Clock Module 1.0 ok/active
B Clock Module 1.0 ok/standby

Fan:

FAN Model Hw Status

Chassis DS-2SLOT-FAN 0.0 ok
PS-1 -- -- ok
PS-2 -- -- absent

Temperature:

Module Sensor MajorThresh MinorThres CurTemp Status
 (Celsius) (Celsius) (Celsius)

1 1 75 60 32 ok
1 2 65 50 32 ok
1 3 -127 -127 43 ok
1 4 -127 -127 39 ok

Power Supply:

PS Model Power Power Status
 (Watts) (Amp @42V)

1 PWR-950-AC 919.38 21.89 ok
2 -- -- -- absent
Mod Model Power Power Power Power Status
 Requested Requested Allocated Allocated
 (Watts) (Amp @42V) (Watts) (Amp @42V)

1 DS-X9216-K9-SUP 220.08 5.24 220.08 5.24 powered-up

Power Usage Summary:

Power Supply redundancy mode: redundant
Total Power Capacity 919.38 W
Power reserved for Supervisor(s) [-] 220.08 W
Power reserved for Fan Module(s) [-] 0.00 W
Power currently used by Modules [-] 0.00 W

Total Power Available 699.30 W

```





## Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs). VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space which allows identical Fibre Channel IDs (FCIDs) to be used simultaneously in different VSANs. VSANs offer the following advantages:

- **Traffic isolation**—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- **Scalability**—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- **Per VSAN fabric services**—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- **Redundancy**—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection is provided (to another VSAN in the same physical SAN) by a configured backup path between the host and the device.
- **Ease of configuration**—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

This chapter includes the following sections:

- [How VSANs Work, page 9-2](#)
- [VSANs Versus Zones, page 9-4](#)
- [Default and Isolated VSANs, page 9-5](#)
- [VSAN Membership, page 9-5](#)
- [VSAN Attributes, page 9-6](#)
- [Creating and Configuring VSANs, page 9-6](#)
- [Assigning VSAN Membership, page 9-7](#)
- [Deleting VSANs, page 9-8](#)
- [Viewing VSAN Configurations, page 9-9](#)
- [Default Settings, page 9-10](#)

# How VSANs Work

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same behavior and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FCIDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

Figure 9-1 shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. Within each VSAN, all members can talk to one another. Between VSANs no communication is possible.

**Figure 9-1 Logical VSAN Segmentation**

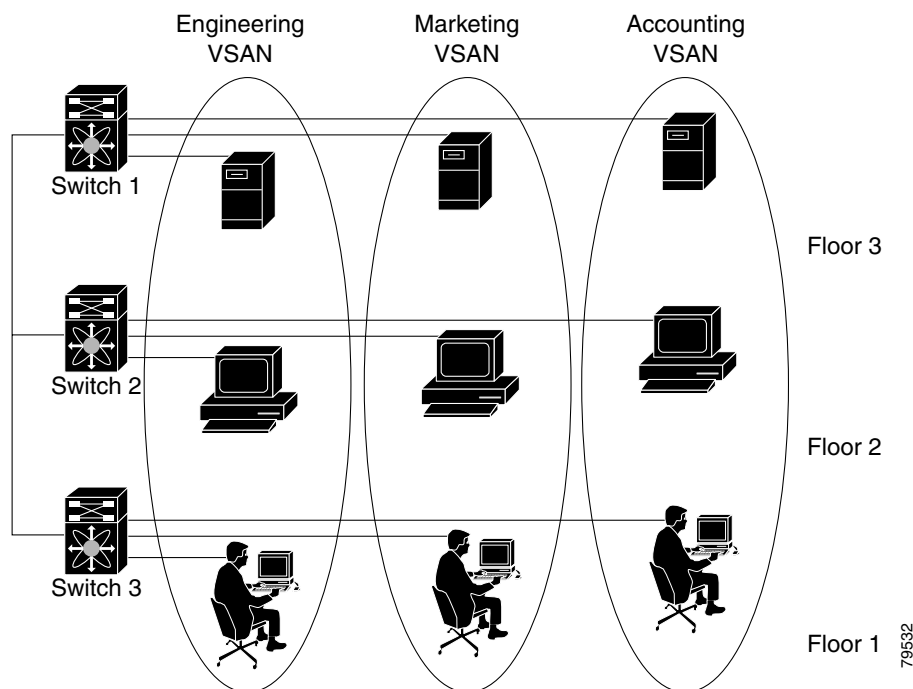
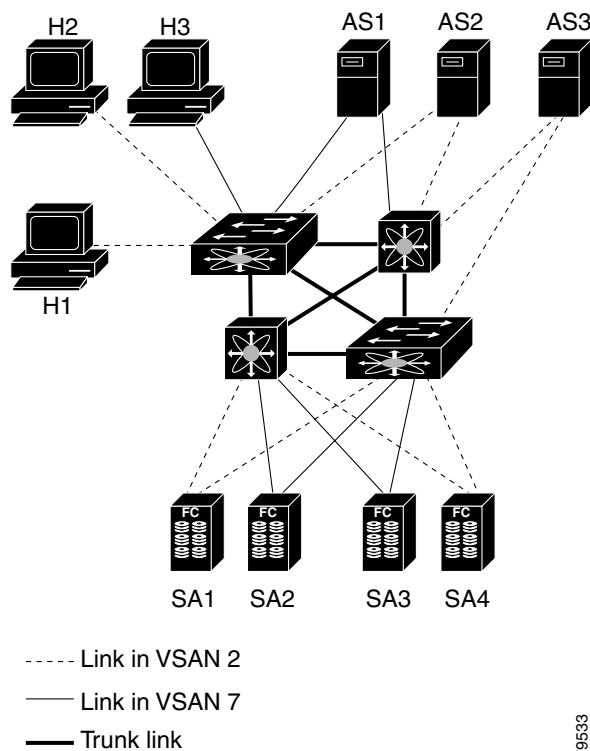


Figure 9-2 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.



As displayed in both [Figure 9-2](#) and [Figure 9-2](#), the switch icons indicate that these features apply to any switch in the Cisco MDS 9000 family.

**Figure 9-2 Example of two VSANs**



The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic. Thus the inter-switch topology of both VSAN 2 and VSAN 7 are identical. This is not a requirement and a network administrator can enable certain VSANs on certain links to create different VSAN topologies.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. [Figure 9-2](#) illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
  - Different customers in storage provider data centers
  - Production or test in an enterprise network
  - Low and high security requirements
  - Backup traffic on separate VSANs
  - Replicating data from user traffic
- VSANs can meet the needs of a particular department or application.

# VSANs Versus Zones

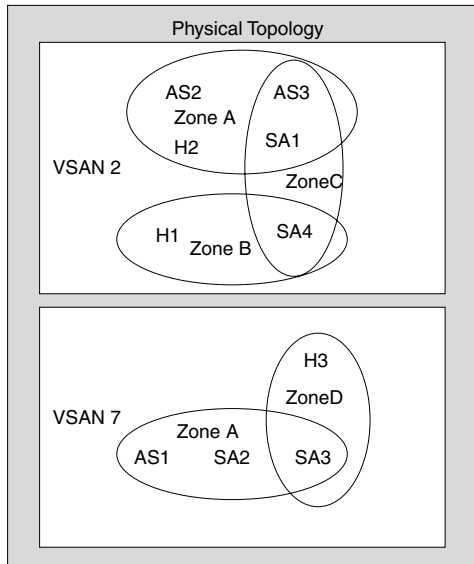
You can define multiple zones in a VSAN. Because two VSANs are equivalent to two nonconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. Table 9-1 lists the differences between VSANs and zones.

Table 9-1 VSAN and Zone Comparison

| VSANs                                                                                             | Zones                                                                               |
|---------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| VSANs equal SANs with routing, naming, and zoning protocols.                                      | These protocols are not available on a per-zone basis.                              |
| —                                                                                                 | Zones are always contained within a VSAN. Zones never span two VSANs.               |
| VSANs limit unicast, multicast, and broadcast traffic.                                            | Zones limit unicast traffic.                                                        |
| Membership is typically defined using the VSAN ID to Fx ports.                                    | Membership is typically defined by the pWWN.                                        |
| An HBA or a storage device may belong only to a single VSAN—the VSAN associated with the Fx port. | An HBA or storage device can belong to multiple zones.                              |
| VSANs enforce membership at each E port, source port, and destination port.                       | Zones enforce membership only at the source and destination ports.                  |
| VSANs are defined for larger environments (storage service providers).                            | Zones are defined for a set of initiators and targets not visible outside the zone. |
| VSANs encompass the entire fabric.                                                                | Zones are configured at the fabric edge.                                            |

Figure 9-3 shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary—they are completely contained within the VSAN. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Figure 9-3 VSANS with Zoning



# Default and Isolated VSANs

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

## Default VSANs

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. If you do not need more than one VSAN for a switch, use this default VSAN as the implicit parameter during configuration. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.

**Note**

VSAN 1 cannot be deleted. It can be suspended.

## Isolated VSANs

VSAN 4094 is an isolated VSAN. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).

**Note**

When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.

**Caution**

Do not use an isolated VSAN to configure ports.

The **show vsan 4094 membership** command displays all ports associated with the isolated VSAN.

## VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis.

By default each port belongs to the default VSAN. You can change the VSAN membership by using the **vsan number interface type port/slot** command.

Trunking ports have an associated list of VSANs that are part of an allowed list (see [Chapter 11, “Configuring Trunking”](#)).

# VSAN Attributes

VSANs have the following attributes:

- **VSAN ID**—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- **State**—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
  - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
  - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN’s configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.
- **VSAN name**—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.



**Note** A VSAN name must be unique.

- **Load balancing attributes**—These attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load balancing path selection.

## Operational State of a VSAN

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

## Creating and Configuring VSANs

You cannot configure any application-specific parameters for a VSAN before creating the VSAN.

To create and configure VSANs, follow these steps:

|        | Command                                                                                         | Purpose                                                                                                         |
|--------|-------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                         | Enters configuration mode.                                                                                      |
| Step 2 | switch(config)# <b>vsan database</b><br>switch(config-vsan-db)#                                 | Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt. |
| Step 3 | switch(config-vsan-db)# <b>vsan 2</b><br>switch(config-vsan-db)#                                | Creates a VSAN with the specified ID (2) if that VSAN does not exist already.                                   |
|        | switch(config-vsan-db)# <b>vsan 2 name TechDoc</b><br>updated vsan 2<br>switch(config-vsan-db)# | Updates the VSAN with the assigned name (TechDoc).                                                              |

|        | Command                                                                                                | Purpose                                                                                                                                                |
|--------|--------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 4 | switch(config-vsan-db)# <b>vsan 2</b><br><b>loadbalancing src-dst-id</b><br>switch(config-vsan-db)#    | Enables the load balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process. |
|        | switch(config-vsan-db)# <b>no vsan 2</b><br><b>loadbalancing src-dst-id</b><br>switch(config-vsan-db)# | Negates the command issued in the previous step and reverts to the default values of the load-balancing parameters.                                    |
|        | switch(config-vsan-db)# <b>vsan 2</b><br><b>loadbalancing src-dst-ox-id</b><br>switch(config-vsan-db)# | Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default).                                                  |
| Step 5 | switch(config-vsan-db)# <b>vsan 2 suspend</b><br>switch(config-vsan-db)#                               | Suspends the selected VSAN.                                                                                                                            |
|        | switch(config-vsan-db)# <b>no vsan 2 suspend</b><br>vs.-config-vsan-db#                                | Negates the <b>suspend</b> command issued in the previous step.                                                                                        |
| Step 6 | switch(config-vsan-db)# <b>end</b><br>switch#                                                          | Returns you to EXEC mode.                                                                                                                              |

## Assigning VSAN Membership

To assign VSAN membership, follow these steps:

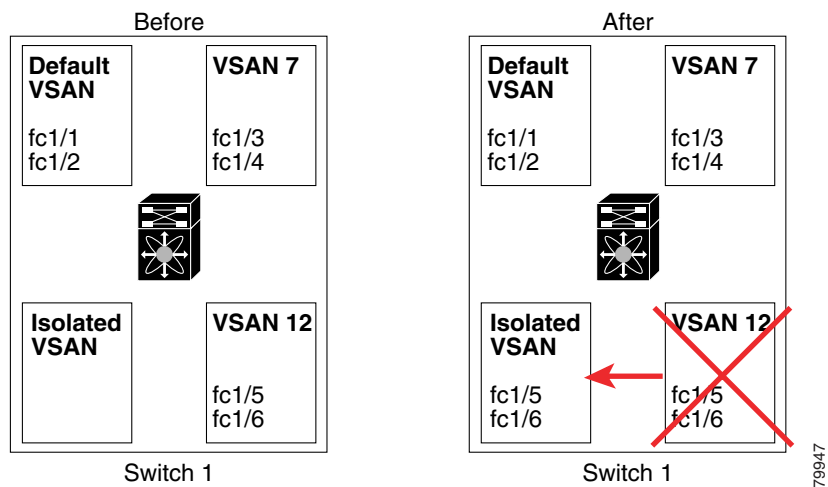
|        | Command                                                                          | Purpose                                                                             |
|--------|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                          | Enters configuration mode.                                                          |
| Step 2 | switch(config)# <b>vsan database</b><br>switch(config-vsan-db)#                  | Configures the database for a VSAN.                                                 |
| Step 3 | switch(config-vsan-db)# <b>vsan 2</b><br>switch(config-vsan-db)#                 | Creates a VSAN with the specified ID (2) if that VSAN does not exist already.       |
| Step 4 | switch(config-vsan-db)# <b>vsan 2 interface fc1/8</b><br>switch(config-vsan-db)# | Assigns the membership of the fc1/8 interface to the specified VSAN (VSAN 2).       |
| Step 5 | switch(config-vsan-db)# <b>vsan 7</b><br>switch(config-vsan-db)#                 | Creates another VSAN with the specified ID (7) if that VSAN does not exist already. |
| Step 6 | switch(config-vsan-db)# <b>vsan 7 interface fc1/8</b><br>switch(config-vsan-db)# | Updates the membership information of the interface to reflect the changed VSAN.    |

# Deleting VSANs

When an active VSAN is deleted, all of its attributes are removed from the running configuration. VSAN related information is maintained by the system software.

- VSAN attributes and port membership details are maintained by VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated, the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see Figure 9-4).

Figure 9-4 VSAN Port Membership Details



- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

  
**Note**

The allowed VSAN list is not affected when a VSAN is deleted (see Chapter 11, “Configuring Trunking”).

Any commands for a nonconfigured VSAN are rejected. For example, if VSAN 10 is not configured in the system, then a command request to move a port to VSAN 10 is rejected.

To delete a VSAN and its various attributes, follow these steps:

|        | Command                                                    | Purpose                                |
|--------|------------------------------------------------------------|----------------------------------------|
| Step 1 | switch# <b>config t</b>                                    | Enters configuration mode.             |
| Step 2 | switch(config)# <b>vsan database</b><br>switch(config-db)# | Configures the VSAN database.          |
| Step 3 | switch-config-db# <b>vsan 2</b><br>switch(config-vsan-db)# | Places you in VSAN configuration mode. |

|        | Command                                                             | Purpose                                      |
|--------|---------------------------------------------------------------------|----------------------------------------------|
| Step 4 | switch(config-vsan-db)# <b>no vsan 5</b><br>switch(config-vsan-db)# | Deletes VSAN 5 from the database and switch. |
| Step 5 | switch(config-vsan-db)# <b>end</b><br>switch#                       | Places you in EXEC mode.                     |

## Viewing VSAN Configurations

Use the **show vsan** command to display information about configured VSANs (see Examples 9-1 to 9-6).

### Example 9-1 Displays the Configuration for a Specific VSAN

```
switch# show vsan 100
vsan 100 information
 name:VSAN0100 state:active
 in-order guarantee:no interoperability mode:no
 loadbalancing:src-id/dst-id/oxid
```

### Example 9-2 Displays the VSAN Usage

```
switch# show vsan usage
4 vsan configured
configured vsans:1-4
vsans available for configuration:5-4093
```

### Example 9-3 Displays All VSANs

```
switch# show vsan
vsan 1 information
 name:VSAN0001 state:active
 in-order guarantee:no interoperability mode:no
 loadbalancing:src-id/dst-id/oxid
vsan 2 information
 name:VSAN0002 state:active
 in-order guarantee:no interoperability mode:no
 loadbalancing:src-id/dst-id/oxid
vsan 7 information
 name:VSAN0007 state:active
 in-order guarantee:no interoperability mode:no
 loadbalancing:src-id/dst-id/oxid
vsan 100 information
 name:VSAN0100 state:active
 in-order guarantee:no interoperability mode:no
 loadbalancing:src-id/dst-id/oxid
vsan 4094:isolated vsan
```

### Example 9-4 Displays Membership Information for the Specified VSAN

```
switch # show vsan 1 membership
vsan 1 interfaces:
 fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 fc1/6 fc1/7 fc1/9
 fc1/10 fc1/11 fc1/12 fc1/13 fc1/14 fc1/15 fc1/16 port-channel 99
```



#### Note

Interface information is not displayed if interfaces are not configured on this VSAN.

**Example 9-5    Displays Membership Information for All VSANs**

```
switch # show vsan membership
vsan 1 interfaces:
 fc2/16 fc2/15 fc2/14 fc2/13 fc2/12 fc2/11 fc2/10 fc2/9
 fc2/8 fc2/7 fc2/6 fc2/5 fc2/4 fc2/3 fc2/2 fc2/1
 fc1/16 fc1/15 fc1/14 fc1/13 fc1/12 fc1/11 fc1/10 fc1/9
 fc1/7 fc1/6 fc1/5 fc1/4 fc1/3 fc1/2 fc1/1
vsan 2 interfaces:
vsan 7 interfaces:
 fc1/8
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

**Example 9-6    Displays Membership Information for a Specified Interface**

```
switch # show vsan membership interface fc1/1
fc1/1
 vsan:1
 allowed list:1-4093
```

# Default Settings

Table 9-2 lists the default settings for all configured VSANs.

**Table 9-2    Default VSAN Parameters**

| Parameters               | Default                                                                                                  |
|--------------------------|----------------------------------------------------------------------------------------------------------|
| State                    | Active state.                                                                                            |
| Name                     | Concatenation of VSAN and a four-digit string representing the VSAN ID. For example, VSAN 3 is VSAN0003. |
| Load-balancing attribute | OX ID (src-dst-ox-id).                                                                                   |
| Port membership          | Default VSAN (VSAN 1).                                                                                   |





## Configuring Interfaces

---

A switch's main function is to relay frames from one data link to another. To do that, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, management interface (mgmt0), or VSAN interfaces.

This chapter describes the basic interface configuration to get your switch up and running. It includes the following sections:

- [Configuring Fibre Channel Interfaces, page 10-2](#)
- [Default Settings, page 10-15](#)
- [Configuring the Management Interface, page 10-16](#)
- [Configuring VSAN Interfaces, page 10-17](#)
- [Displaying Interface Information, page 10-17](#)



### Note

See [Chapter 4, “Initial Configuration”](#) and [Chapter 20, “Configuring IP Services,”](#) for more information on configuring mgmt0 interfaces.

Before you begin configuring the switch, ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode (see the [“Verifying the Module Status”](#) section on page 4-15).

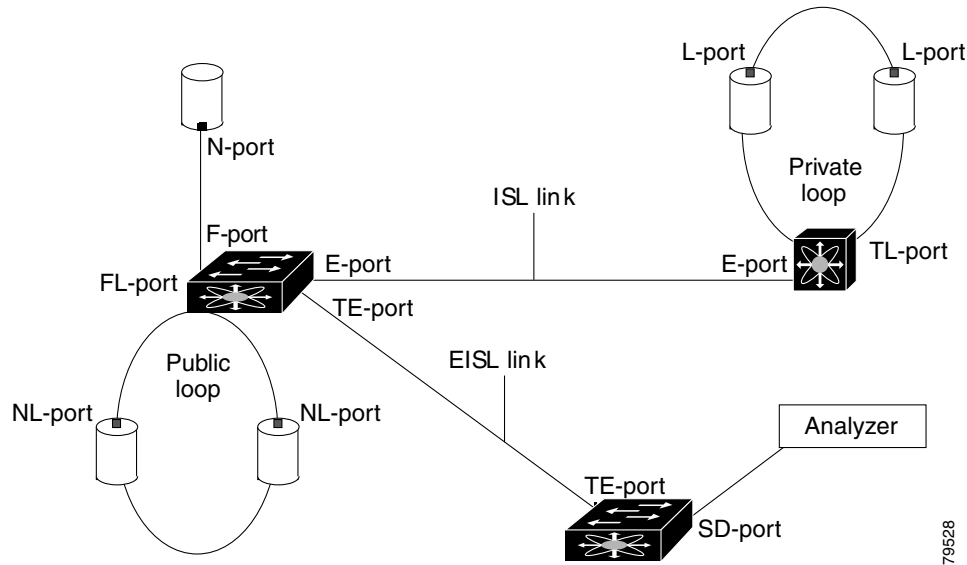
# Configuring Fibre Channel Interfaces

This section describes Fibre Channel interface characteristics, including (but are not limited to) modes, states, and speeds. It includes the following sections:

- [About Interface Modes, page 10-2](#)
- [About Interface States, page 10-6](#)
- [Configuring Fibre Channel Interfaces, page 10-9](#)
- [Configuring a Range of Interfaces, page 10-9](#)
- [Disabling Interfaces, page 10-9](#)
- [Configuring Interface Modes, page 10-9](#)
- [Configuring Administrative Speeds, page 10-10](#)
- [Configuring Interface Descriptions, page 10-10](#)
- [Configuring Buffer-to-Buffer Credits, page 10-11](#)
- [Configuring Performance Buffers, page 10-12](#)
- [Configuring Frame Encapsulation, page 10-12](#)
- [Configuring Receive Data Field Size, page 10-12](#)
- [Configuring the Beacon Mode, page 10-13](#)
- [Identifying the Beacon LEDs, page 10-13](#)
- [Configuring Switch Port Defaults, page 10-14](#)
- [Identifying FCOT Transmitter Types, page 10-14](#)

## About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several modes: E port, F port, FL port, TL port, TE port, and SD port (see [Figure 10-1](#)). Besides these modes, each interface may be configured in auto or Fx port mode. These two modes determine the port type during interface initialization. A brief description of each interface mode follows.

**Figure 10-1 Cisco MDS 9000 Family Switch Interface Modes****Note**

Interfaces are created in VSAN 1 by default. See [Chapter 9, “Configuring and Managing VSANs.”](#)

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

A brief description of each interface mode follows.

## E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see [Chapter 12, “Configuring PortChannels”](#)).

## F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

## FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

## TL Port

In translatable loop port (TL port) mode, an interface functions as a translatable loop port. It may be connected to one or more private loop devices (NL ports). TL port mode is specific to Cisco MDS 9000 family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

See the [“Displaying TL Port Information” section on page 10-27](#). TL ports support class 2 and class 3 services.

**Note**

---

Devices attached to TL ports are recommended to be configured in zones which have up to 64 zone members.

---

## TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an Extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (**fctrace**) feature

In TE-port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family (see [Chapter 11, “Configuring Trunking”](#)). TE ports support class 2, class 3, and class F service.

## SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes through a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, they

merely transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see [Chapter 28, “Monitoring Network Traffic Using SPAN”](#)).

## ST Port

In the SPAN Tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus, cannot be used for normal Fibre Channel traffic.

(see the [“Remote SPAN”](#) section on page 28-15).

## Fx Port

Interfaces configured as Fx ports are allowed to operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

## B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2. [Figure 22-11](#) depicts a typical SAN extension over an IP network.

When an FCIP peer is a SAN extender device that only support Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see [Chapter 22, “Configuring IP Storage”](#)).

## Auto Mode

Interfaces configured as **auto** are allowed to operate in one of the following modes: F port, FL port, E port, or TE port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode (see [Chapter 11, “Configuring Trunking”](#)). TL ports and SD ports are not determined during initialization and are administratively configured.

## About Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

### Administrative States

The administrative state refers to the administrative configuration of the interface as described in [Table 10-1](#).

**Table 10-1 Administrative States**

| Administrative State | Description                                                                                                                                        |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Up                   | Enables an interface.                                                                                                                              |
| Down                 | Disables an interface. When an interface is administratively disabled ( <b>shutdown</b> command), the physical link layer state change is ignored. |

### Operational States

The operational state indicates the current operational state of the interface as described in [Table 10-2](#).

**Table 10-2 Operational States**

| Operational State | Description                                                                                                                                                                                                              |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Up                | Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed. |
| Down              | Interface cannot transmit or receive (data) traffic.                                                                                                                                                                     |
| Trunking          | Interface is operational in TE mode.                                                                                                                                                                                     |

### Reason Codes

Reason codes are dependent on the operational state of the interface as described in [Table 10-3](#).

**Table 10-3 Reason Codes for Interface States**

| Administrative Configuration | Operational Status | Reason Code                                                                                                                                     |
|------------------------------|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Up                           | Up                 | None.                                                                                                                                           |
| Down                         | Down               | Administratively down—If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted. |
| Up                           | Down               | See <a href="#">Table 10-4</a> .                                                                                                                |

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in [Table 10-4](#).

**Table 10-4 Reason Codes for Nonoperational States**

| Reason Code                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Applicable Modes          |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
| Link failure or not connected                   | Physical layer link is not operational.                                                                                                                                                                                                                                                                                                                                                                                                                           | All                       |
| Fcot not present                                | The Fibre Channel optical transmitter hardware (FCOT) is not plugged in.                                                                                                                                                                                                                                                                                                                                                                                          |                           |
| Initializing                                    | The physical layer link is operational and the protocol initialization is in progress.                                                                                                                                                                                                                                                                                                                                                                            |                           |
| Reconfigure fabric in progress                  | The fabric is currently being reconfigured.                                                                                                                                                                                                                                                                                                                                                                                                                       |                           |
| Offline                                         | Waiting for the specified R_A_TOV time before retrying initialization.                                                                                                                                                                                                                                                                                                                                                                                            |                           |
| Inactive                                        | The interface VSAN is deleted or is in a suspended state.<br><br>To make the interface operational, assign that port to a configured and active VSAN.                                                                                                                                                                                                                                                                                                             |                           |
| Hardware failure                                | A hardware failure is detected.                                                                                                                                                                                                                                                                                                                                                                                                                                   |                           |
| Error disabled                                  | Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: <ul style="list-style-type: none"> <li>Configuration failure.</li> <li>Incompatible buffer-to-buffer credit configuration.</li> </ul> To make the interface operational, you must first fix the error conditions causing this state; and next, administratively configure the interface as <b>shutdown</b> followed by <b>no shutdown</b> . |                           |
| Isolation due to ELP failure                    | Port negotiation failed.                                                                                                                                                                                                                                                                                                                                                                                                                                          | Only E ports and TE ports |
| Isolation due to ESC failure                    | Port negotiation failed.                                                                                                                                                                                                                                                                                                                                                                                                                                          |                           |
| Isolation due to domain overlap                 | The Fibre Channel domains (fcdomain) overlap.                                                                                                                                                                                                                                                                                                                                                                                                                     |                           |
| Isolation due to domain ID assignment failure   | The assigned domain ID is not valid.                                                                                                                                                                                                                                                                                                                                                                                                                              |                           |
| Isolation due to other side E port isolated     | The E port at the other end of the link is isolated.                                                                                                                                                                                                                                                                                                                                                                                                              |                           |
| Isolation due to invalid fabric reconfiguration | The port is isolated due to fabric reconfiguration.                                                                                                                                                                                                                                                                                                                                                                                                               |                           |
| Isolation due to domain manager disabled        | The fcdomain feature is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                 |                           |
| Isolation due to zone merge failure             | The zone merge operation failed.                                                                                                                                                                                                                                                                                                                                                                                                                                  |                           |
| Isolation due to VSAN mismatch                  | The VSANs at both ends of an ISL are different.                                                                                                                                                                                                                                                                                                                                                                                                                   |                           |

**Table 10-4 Reason Codes for Nonoperational States (continued)**

| Reason Code                                     | Description                                                                                                                                                                                                | Applicable Modes            |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| Nonparticipating                                | FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode. | Only FL ports and TL ports  |
| PortChannel administratively down               | The interfaces belonging to the PortChannel are down.                                                                                                                                                      | Only PortChannel interfaces |
| Suspended due to incompatible speed             | The interfaces belonging to the PortChannel have incompatible speeds.                                                                                                                                      |                             |
| Suspended due to incompatible mode              | The interfaces belonging to the PortChannel have incompatible modes.                                                                                                                                       |                             |
| Suspended due to incompatible remote switch WWN | An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.                                                                                        |                             |

### Configuring 32-port Switching Modules and Host-Optimized Ports

The 32-port 1/2-Gbps switching module contains 8 port groups of 4 ports each. When configuring these modules or the host-optimized ports in the Cisco 9100 Series, the following guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain in the shutdown state.
- If any of the other three ports are configured in a no shutdown state, you cannot configure the first port as an E port. The other three ports continue to remain in a no shutdown state.
- The default port mode is **auto**. The **auto** option is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules and the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).



#### Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.



## Configuring Fibre Channel Interfaces

To configure a Fibre Channel interface, follow these steps:

|        | Command                                | Purpose                             |
|--------|----------------------------------------|-------------------------------------|
| Step 1 | switch# <b>config t</b>                | Enters configuration mode.          |
| Step 2 | switch(config)# <b>interface fc1/1</b> | Configures the specified interface. |

When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID).

## Configuring a Range of Interfaces

To configure a range of interfaces, follow these steps:

|        | Command                                                | Purpose                                                                  |
|--------|--------------------------------------------------------|--------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                | Enters configuration mode.                                               |
| Step 2 | switch(config)# <b>interface fc1/1 - 4 , fc2/1 - 3</b> | Configures the range of specified interfaces.                            |
|        |                                                        | <b>Note</b> In this command, provide a space before and after the comma. |

## Disabling Interfaces

Interfaces on a port are shut down by default (unless you modified the initial configuration). To enable traffic flow, follow these steps:

|        | Command                                | Purpose                                                                                                                          |
|--------|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                | Enters configuration mode.                                                                                                       |
| Step 2 | switch(config)# <b>interface fc1/1</b> | Configures the specified interface.                                                                                              |
| Step 3 | switch(config-if)# <b>no shutdown</b>  | Enables traffic flow to administratively allow traffic when the <b>no</b> prefix is used (provided the operational state is up). |
|        | switch(config-if)# <b>shutdown</b>     | Shuts down the interface and disables traffic flow (default).                                                                    |

## Configuring Interface Modes

To configure the interface mode, follow these steps:

|        | Command                                                      | Purpose                             |
|--------|--------------------------------------------------------------|-------------------------------------|
| Step 1 | switch# <b>config t</b>                                      | Enters configuration mode.          |
| Step 2 | switch(config)# <b>interface fc1/1</b><br>switch(config-if)# | Configures the specified interface. |

|        | Command                                                              | Purpose                                                                                                                                                                                                                                     |
|--------|----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 3 | switch(config-if)# <b>switchport mode F</b><br>switch(config-if)#    | Configures the administrative mode of the port. You can set the operational state to auto, E, F, FL, Fx, TL, or SD port mode.<br><br><b>Note</b> Fx ports refers to an F port or an FL port (host connection only), but not E ports.        |
|        | switch(config-if)# <b>switchport mode auto</b><br>switch(config-if)# | Configures the interface mode to auto-negotiate an E, F, FL, or TE port mode (not TL or SD-port modes) of operation.<br><br><b>Note</b> TL ports and SD ports cannot be configured automatically. They must be administratively configured. |

## Configuring Administrative Speeds

By default, the administrative speed for an interface is automatically calculated by the switch. To configure the administrative speed of the interface, follow these steps:

|        | Command                                                               | Purpose                                                                                                                                                                                                                                                  |
|--------|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                               | Enters configuration mode.                                                                                                                                                                                                                               |
| Step 2 | switch(config-if)# <b>switchport speed 1000</b><br>switch(config-if)# | Configures the administrative speed of the interface to 1000 Mbps.<br><br>The number indicates the speed in megabits per second (Mbps). You can set the speed to 1000 Mbps (for 1Gbps interfaces), 2000 Mbps (for 2 Gbps interfaces), or auto (default). |
|        | switch(config-if)# <b>switchport speed auto</b><br>switch(config-if)# | Reconfigures the factory default (auto) administrative speed of the interface.                                                                                                                                                                           |

## Configuring Interface Descriptions

To configure a description for an interface, follow these steps:

|        | Command                                                      | Purpose                                                                                         |
|--------|--------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                      | Enters configuration mode.                                                                      |
| Step 2 | switch(config)# <b>interface fc1/1</b><br>switch(config-if)# | Configures the specified interface.                                                             |
| Step 3 | switch(config-if)# <b>switchport description cisco-HBA2</b>  | Configures the description of the interface.<br><br>The string may be up to 80 characters long. |
|        | switch(config-if)# <b>no switchport description</b>          | Clears the description of the interface.                                                        |

## Configuring Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB\_credits) are a flow control mechanism to ensure that FC switches do not run out of buffers, since switches must not drop frames. Buffer Credits are negotiated on a per-hop basis.

The receive BB\_credit (rxbbcredit) value may be configured for each FC interface. In most cases, you don't need to modify the default configuration.



### Note

The receive BB\_credit values depend on the module type and the port mode:

16-port switching modules and full rate ports: The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.

32-port switching modules and host-optimized ports: The default value is 12 for the Fx, E, and TE modes. These values cannot be changed.

To configure buffer-to-buffer credits for a Fibre Channel interface, follow these steps:

|        | Command                                                                                                                      | Purpose                                                                                                                                                                                                                                                                      |
|--------|------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                      | Enters configuration mode.                                                                                                                                                                                                                                                   |
| Step 2 | switch(config)# <b>interface fc1/1</b><br>switch(config-if)#                                                                 | Configures the specified interface.                                                                                                                                                                                                                                          |
| Step 3 | switch(config-if)# <b>switchport fcrxbbcredit default</b>                                                                    | Applies the default operational value to the selected interface. The operational value depends on the port mode. The default values are assigned based on the port capabilities.                                                                                             |
|        | switch(config-if)# <b>switchport fcrxbbcredit 5</b>                                                                          | Assigns a BB_credit of 5 to the selected interface. The range to assign BB_credits is between 1 and 255.                                                                                                                                                                     |
|        | switch(config-if)# <b>switchport fcrxbbcredit 5 mode E</b>                                                                   | Assigns this value if the port is operating in E or TE mode. The range to assign BB_credits is between 1 and 255.                                                                                                                                                            |
|        | switch(config-if)# <b>switchport fcrxbbcredit 5 mode Fx</b>                                                                  | Assigns this value if the port is operating in F or FL mode. The range to assign BB_credits is between 1 and 255.                                                                                                                                                            |
| Step 4 | switch# <b>do show int fc1/1</b><br>fc1/1 is up<br>...<br>16 receive B2B credit remaining<br>3 transmit B2B credit remaining | Displays the receive and transmit BB_credit along with other pertinent interface information for this interface.<br><br><b>Note</b> The BB_credit values are correct at the time the registers are read. They are useful to verify situations when the data traffic is slow. |



### Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

## Configuring Performance Buffers

Regardless of the configured Rx BB\_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB\_credit value.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used.

If you do not specify this command, the **default** option is automatically used.

To configure performance buffers for a Fibre Channel interface, follow these steps:

|        | Command                                                                      | Purpose                                                                                                                                                               |
|--------|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                      | Enters configuration mode.                                                                                                                                            |
| Step 2 | switch(config)# <b>interface fc1/1</b><br>switch(config-if)#                 | Configures the specified interface.                                                                                                                                   |
| Step 3 | switch(config-if)# <b>switchport fcrxbcredit performance-buffers 45</b>      | Assigns a performance buffer of 45 to the selected interface. The performance buffer value ranges from 1 and 145.                                                     |
|        | switch(config-if)# <b>switchport fcrxbcredit performance-buffers default</b> | Reverts to the factory default of using the built-in algorithm.<br><br>The <b>show interface</b> command displays the performance buffer value if default is changed. |

## Configuring Frame Encapsulation

The **switchport encap eisl** command only applies to SD port interfaces. This command determines the frame format for all frames transmitted by the interface in SD port mode. When the encap is set to EISL, all frames are transmitted in the EISL frame format irrespective of the SPAN source(s).

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames will be encapsulated, and you will see a new line (Encapsulation is eisl) in the **show interface SD\_port\_interface** command output (see the “[Encapsulating Frames](#)” section on page 28-8).

## Configuring Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces by issuing the **switchport fcrxbuFSIZE** command. The default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure data field size for a Fibre Channel interface, follow these steps:

|        | Command                 | Purpose                    |
|--------|-------------------------|----------------------------|
| Step 1 | switch# <b>config t</b> | Enters configuration mode. |

|        | Command                                                      | Purpose                                                                                                                                  |
|--------|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Step 2 | switch(config)# <b>interface</b> fc1/1<br>switch(config-if)# | Configures the specified interface.                                                                                                      |
| Step 3 | switch(config-if)# <b>switchport fcrxbufsize</b> 2000        | Reduces the data field size for the selected interface to 2000 bytes. The default is 2112 bytes and the range is from 256 to 2112 bytes. |

## Configuring the Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface. The **beacon** command has no effect on the operation of the interface.

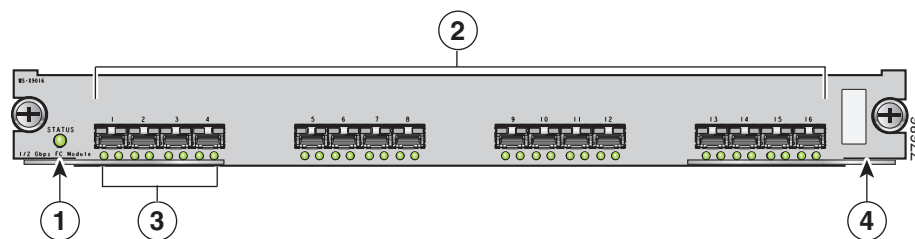
To enable beacon mode for a specified interface or range of interfaces, follow these steps:

|        | Command                                                                                           | Purpose                                                                                       |
|--------|---------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                                        | Enters configuration mode.                                                                    |
| Step 2 | switch(config)# <b>interface</b> fc1/1<br>switch(config-if)#                                      | Configures the specified interface.                                                           |
| Step 3 | switch(config-if)# <b>switchport beacon</b><br><br>switch(config-if)# <b>no switchport beacon</b> | Enables the beacon mode for the interface.<br><br>Disables the beacon mode for the interface. |

## Identifying the Beacon LEDs

Figure 10-2 displays the status, link, and speed LEDs in a 16-port switching module.

**Figure 10-2 Cisco MDS 9000 Family Switch Interface Modes**



|   |                                                                                                                                   |   |                                                                                                               |
|---|-----------------------------------------------------------------------------------------------------------------------------------|---|---------------------------------------------------------------------------------------------------------------|
| 1 | Status LED (see the “Identifying Module LEDs” section on page 7-8)                                                                | 3 | Link LEDs (see the “Identifying Module LEDs” section on page 7-8) and speed LEDs (explained in this section). |
| 2 | 1/2-Gbps Fibre Channel port group (see the “Configuring 32-port Switching Modules and Host-Optimized Ports” section on page 10-8) | 4 | Asset tag (refer to the <i>Cisco MDS 9000 Family Hardware Installation Guide</i> ).                           |

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—the interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—the interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off—beacon mode is disabled
- On (flashing green)—the beacon mode is enabled. The LED flashes at one-second intervals.

## Configuring Switch Port Defaults

You can configure default values for various switch port attributes. If you configure the following attributes, they will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

To configure switch port attributes, follow these steps:

|        | Command                                                                                | Purpose                                                                                                                                                                                                                                                    |
|--------|----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                | Enters configuration mode.                                                                                                                                                                                                                                 |
| Step 2 | switch(config)# <b>no system default switchport shutdown</b><br>switch(config-if)#     | Configures the default setting for administrative state of an interface as Up. (The factory default setting is Down).<br><br><b>Tip</b> This command is applicable only to interfaces for which no user configuration exists for the administrative state. |
|        | switch(config)# <b>system default switchport shutdown</b><br>switch(config-if)#        | Configures the default setting for administrative state of an interface as Down. This is the factory default setting.<br><br><b>Tip</b> This command is applicable only to interfaces for which no user configuration exists for the administrative state. |
|        | switch(config)# <b>system default switchport trunk mode auto</b><br>switch(config-if)# | Configures the default setting for administrative trunk mode state of an interface as Auto.<br>(The factory default setting is trunk mode On).                                                                                                             |

## Identifying FCOT Transmitter Types

The FCOT transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related FCOT has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fcslot/port transceiver** command display both values for Cisco supported FCOTs (see the [“Displaying Interface Information”](#) section on page 10-17).

**Table 10-5 FCOT Trimeter Acronym Definitions**

| Definition                                                      | Acronym |
|-----------------------------------------------------------------|---------|
| <b>Standard transmitters defined in the GBIC Specifications</b> |         |
| short wave laser                                                | swl     |
| long wave laser                                                 | lwl     |
| long wave laser cost reduced                                    | lwcr    |
| electrical                                                      | elec    |
| <b>Extended transmitters assigned to Cisco-supported FCOTs</b>  |         |
| CWDM-1470                                                       | c1470   |
| CWDM-1490                                                       | c1490   |
| CWDM-1510                                                       | c1510   |
| CWDM-1530                                                       | c1530   |
| CWDM-1550                                                       | c1550   |
| CWDM-1570                                                       | c1570   |
| CWDM-1590                                                       | c1590   |
| CWDM-1610                                                       | c1610   |

## Default Settings

Table 10-6 lists the default settings for Fibre Channel interface parameters.

**Table 10-6 Default Fibre Channel Interface Parameters**

| Parameters           | Default                                        |
|----------------------|------------------------------------------------|
| Interface mode       | Auto                                           |
| Interface speed      | Auto                                           |
| Administrative state | Shutdown (unless changed during initial setup) |
| Trunk mode           | On (unless changed during initial setup)       |
| Trunk-allowed VSANs  | 1 to 4093                                      |
| Interface VSAN       | Default VSAN (1)                               |
| Beacon mode          | Off (disabled)                                 |
| EISL encapsulation   | Disabled                                       |
| Data field size      | 2112 bytes                                     |

# Configuring the Management Interface

You can remotely configure the switch through the management interface (mgmt0). To configure a connection remotely, you must configure the IP parameters (IP address, subnet mask, and default gateway) from the CLI so that the switch is reachable.



## Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask.

To configure the mgmt0 Ethernet interface, follow these steps:

|        | Command                                                              | Purpose                                                                                                                                                                        |
|--------|----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config terminal</b><br>switch(config)#                    | Enters configuration mode.                                                                                                                                                     |
| Step 2 | switch(config)# <b>interface mgmt0</b><br>switch(config-if)#         | Configures the management Ethernet interface on the switch to configure the management interface.                                                                              |
| Step 3 | switch(config-if)# <b>ip address 172.16.1.2 255.255.0</b>            | Enters the IP address and IP subnet mask for the interface specified in Step 2.                                                                                                |
| Step 4 | switch(config-if)# <b>no shutdown</b>                                | Enables the interface.                                                                                                                                                         |
| Step 5 | switch(config-if)# <b>exit</b><br>switch(config)#                    | Returns to configuration mode.                                                                                                                                                 |
| Step 6 | switch(config)# <b>ip default-gateway 1.1.1.4</b><br>switch(config)# | Configures the default gateway IP address.                                                                                                                                     |
| Step 7 | switch(config)# <b>exit</b><br>switch#                               | Returns to EXEC mode.                                                                                                                                                          |
| Step 8 | switch# <b>copy running-config startup-config</b>                    | Saves your configuration changes to the file system.<br><br><b>Note</b> This step is optional. If you wish to save your configuration, you can issue this command at any time. |

The management port (mgmt0) is autosensing and operates as full duplex mode and 100 Mbps speed. The speed and mode cannot be configured.



## Note

You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.



## Configuring VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexistent VSANs.

Follow these guidelines when creating or deleting VSAN interfaces:

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface using the **interface VSAN** command. This is not done automatically.
- If you delete the VSAN, the attached interface is automatically deleted.

To create a VSAN interface, follow these steps:

|        | Command                                                       | Purpose                          |
|--------|---------------------------------------------------------------|----------------------------------|
| Step 1 | switch# <b>config t</b>                                       | Enters configuration mode.       |
| Step 2 | switch(config)# <b>interface vsan 5</b><br>switch(config-if)# | Configures a VSAN with the ID 5. |

You can configure each interface only in one VSAN.

After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) features (see [Chapter 20, “Configuring IP Services”](#)).

## Displaying Interface Information

The **show interface** command is invoked from the EXEC mode and displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch. See Examples 10-1 to 10-11.

### Example 10-1 Displays All Interfaces

```
switch# show interface
switch# do show int fc1/1
fc1/1 is up
 Hardware is Fibre Channel, FCOT is short wave laser
 Port WWN is 20:0b:00:05:30:00:8d:de
 Admin port mode is F
 Port mode is F, FCID is 0x610000
 Port vsan is 2
 Speed is 2 Gbps
 Transmit B2B Credit is 3
 Receive B2B Credit is 16
 Receive data field Size is 2112
 Beacon is turned off
 5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 134 frames input, 8468 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
 154 frames output, 46072 bytes
 0 discards, 0 errors
```

```

1 input OLS, 1 LRR, 0 NOS, 0 loop inits
1 output OLS, 0 LRR, 1 NOS, 0 loop inits
16 receive B2B credit remaining
3 transmit B2B credit remaining.
. . .
fc1/9 is trunking
Hardware is Fibre Channel, FCOT is long wave laser cost reduced
Port WWN is 20:09:00:05:30:00:97:9e
Peer port WWN is 20:0b:00:0b:5f:a3:cc:00
Admin port mode is E, trunk mode is on
Port mode is TE
Port vsan is 100
Speed is 2 Gbps
Transmit B2B Credit is 255
Receive B2B Credit is 255
Receive data field Size is 2112
Beacon is turned off
Trunk vsans (admin allowed and active) (1,100,3000)
Trunk vsans (up) (1,100,3000)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
5 minutes input rate 280 bits/sec, 35 bytes/sec, 0 frames/sec
5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
4609939 frames input, 8149405708 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
4638491 frames output, 7264731728 bytes
 0 discards, 0 errors
3 input OLS, 9 LRR, 1 NOS, 0 loop inits
9 output OLS, 7 LRR, 1 NOS, 0 loop inits
16 receive B2B credit remaining
3 transmit B2B credit remaining.
. . .
fc1/13 is up
Hardware is Fibre Channel, FCOT is short wave laser
Port WWN is 20:0d:00:05:30:00:97:9e
Admin port mode is auto, trunk mode is on
Port mode is F, FCID is 0x650100
Port vsan is 100
Speed is 2 Gbps
Transmit B2B Credit is 3
Receive B2B Credit is 16
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
8696 frames input, 3227212 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
16799 frames output, 6782444 bytes
 0 discards, 0 errors
0 input OLS, 0 LRR, 0 NOS, 0 loop inits
1 output OLS, 1 LRR, 0 NOS, 1 loop inits
16 receive B2B credit remaining
3 transmit B2B credit remaining.
. . .
sup-fc0 is up
Hardware is Fibre Channel
Speed is 1 Gbps
139597 packets input, 13852970 bytes
 0 multicast frames, 0 compressed
 0 input errors, 0 frame, 0 overrun 0 fifo

```

```

139516 packets output, 16759004 bytes, 0 underruns
 0 output errors, 0 collisions, 0 fifo
 0 carrier errors

mgmt0 is up
 Hardware is FastEthernet
 Address is 0005.3000.80fe
 Internet address is 172.19.48.96/25
 MTU 1500 bytes, BW 100 Mbps
 321561 packets input, 70215667 bytes
 0 multicast frames, 0 compressed
 0 input errors, 0 frame, 0 overrun 0 fifo
 334550 packets output, 307482596 bytes, 0 underruns
 0 output errors, 0 collisions, 0 fifo
 0 carrier errors

vsan1 is up, line protocol is up
 WWPN is 10:00:00:05:30:00:12:63, FCID is 0xef001e
 Internet address is 10.10.11.10/24
 MTU 1500 bytes, BW 1000000 Kbit
 0 packets input, 0 bytes, 0 errors, 0 multicast
 0 packets output, 0 bytes, 0 errors, 0 dropped
. . .
port-channel 1 is trunking
 Hardware is Fibre Channel
 Port WWN is 24:01:00:05:30:00:97:9e
 Admin port mode is E, trunk mode is on
 Port mode is TE
 Port vsan is 1
 Speed is 4 Gbps
 Trunk vsans (admin allowed and active) (1,100,3000)
 Trunk vsans (up) (1)
 Trunk vsans (isolated) (100,3000)
 Trunk vsans (initializing) ()
 5 minutes input rate 648 bits/sec, 81 bytes/sec, 0 frames/sec
 5 minutes output rate 304 bits/sec, 38 bytes/sec, 0 frames/sec
 4629945 frames input, 206672020 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
 4547515 frames output, 687414748 bytes
 0 discards, 0 errors
 2 input OLS, 2 LRR, 4 NOS, 0 loop inits
 6 output OLS, 2 LRR, 4 NOS, 0 loop inits
 Member[1] : fc1/1
 Member[2] : fc1/2.
. . .

```

You can also specify arguments (a range of interfaces or multiple, specified interfaces) to display interface information. You can specify a range of interfaces by issuing a command with the following example format:

```
interface fc1/1 - 5 , fc2/5 - 7
```


**Note**

The spaces are required before and after the dash ( - ) and before and after the comma ( , )

**Example 10-2 Displays Multiple, Specified Interfaces**

```

switch# show interface fc3/13 , fc3/16
fc3/13 is up
 Hardware is Fibre Channel, FCOT is short wave laser
 Port WWN is 20:8d:00:05:30:00:97:9e

```

```

Admin port mode is FX
Port mode is F, FCID is 0x7b0300
Port vsan is 1
Speed is 2 Gbps
Transmit B2B Credit is 3
Receive B2B Credit is 12
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 1856 frames input, 116632 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
 1886 frames output, 887712 bytes
 0 discards, 0 errors
 0 input OLS, 0 LRR, 0 NOS, 1 loop inits
 1 output OLS, 1 LRR, 0 NOS, 1 loop inits
 16 receive B2B credit remaining
 3 transmit B2B credit remaining.

fc3/16 is up
Hardware is Fibre Channel, FCOT is short wave laser
Port WWN is 20:90:00:05:30:00:97:9e
Admin port mode is FX
Port mode is F, FCID is 0x7d0100
Port vsan is 3000
Speed is 2 Gbps
Transmit B2B Credit is 3
Receive B2B Credit is 12
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 504 bits/sec, 63 bytes/sec, 0 frames/sec
5 minutes output rate 520 bits/sec, 65 bytes/sec, 0 frames/sec
 47050 frames input, 10311824 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
 62659 frames output, 10676988 bytes
 0 discards, 0 errors
 0 input OLS, 0 LRR, 0 NOS, 0 loop inits
 1 output OLS, 1 LRR, 0 NOS, 1 loop inits
 16 receive B2B credit remaining
 3 transmit B2B credit remaining.

```

### **Example 10-3** *Displays a Specific Interface*

```

switch# show interface fc2/2
fc2/2 is trunking
 Port description is Trunk to Core-4
 Hardware is Fibre Channel, FCOT is short wave laser
 Port WWN is 20:42:00:05:30:00:97:9e
 Peer port WWN is 20:cc:00:05:30:00:50:9e
 Admin port mode is E, trunk mode is on
 Port mode is TE
 Port vsan is 1
 Speed is 2 Gbps
 Transmit B2B Credit is 255
 Receive B2B Credit is 255
 Receive data field Size is 2112
 Beacon is turned off
 Belongs to port-channel 2
 Trunk vsans (admin allowed and active) (1,100,3000)

```

```

Trunk vsans (up) (1)
Trunk vsans (isolated) (100,3000)
Trunk vsans (initializing) ()
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
2214834 frames input, 98673588 bytes
0 discards, 0 errors
0 CRC, 0 unknown class
0 too long, 0 too short
2262415 frames output, 343158368 bytes
0 discards, 0 errors
1 input OLS, 1 LRR, 1 NOS, 0 loop inits
2 output OLS, 1 LRR, 0 NOS, 0 loop inits
16 receive B2B credit remaining
3 transmit B2B credit remaining.

```

#### Example 10-4 Displays a VSAN Interface

```

switch# show interface vsan 2
vsan2 is up, line protocol is up
WWPN is 10:00:00:05:30:00:59:1f, FCID is 0xb90100
Internet address is 10.1.1.1/24
MTU 1500 bytes, BW 1000000 Kbit
0 packets input, 0 bytes, 0 errors, 0 multicast
0 packets output, 0 bytes, 0 errors, 0 dropped

```

#### Example 10-5 Displays Port Description

```
switch# show interface description
```

```

Interface Description

fc3/1 test intest
fc3/2 --
fc3/3 --
fc3/4 TE port
fc3/5 --
fc3/6 --
fc3/10 Next hop switch 5
fc3/11 --
fc3/12 --
fc3/16 --

```

```

Interface Description

port-channel 1 --
port-channel 5 --
port-channel 6 --

```

#### Example 10-6 Displays Interface Information in a Brief Format

```
switch# show interface brief
```

```

Interface Vsan Admin Admin Status FCOT Oper Oper Port
 Mode Trunk Mode Mode Speed Channel
 (Gbps)

fc1/1 1 E on trunking sw1 TE 2 1
fc1/2 1 E on trunking sw1 TE 2 1
fc1/3 1 auto on fcotAbsent -- -- -- --
fc1/4 1 auto on fcotAbsent -- -- -- --

```

```

fc1/5 3000 auto on up swl F 2 --
...
fc2/2 1 E on trunking swl TE 2 2
fc2/3 1 auto on down cl610 -- --
fc2/4 1 auto on down cl590 -- --
fc2/5 3000 auto on notConnected lwcr -- --
fc2/6 1 auto on fcotAbsent -- -- --
...
fc3/16 3000 FX -- up swl F 2 --
fc3/17 1 FX -- fcotAbsent -- -- --
...

```

```

Interface Status IP Address Speed MTU

GigabitEthernet4/1 fcotAbsent -- auto 1500
...
GigabitEthernet4/6 down 10.1.1.2/8 auto 3000
GigabitEthernet4/7 down 10.1.1.27/24 auto 1500
GigabitEthernet4/8 down -- auto 1500

```

```

Interface Status Oper Mode Oper Speed
 (Gbps)

```

```

iscsi4/1 down --
...

```

```

Interface Status Speed
 (Gbps)

```

```

sup-fc0 up 1

```

```

Interface Status IP Address Speed MTU

mgmt0 up 172.19.48.96/25 100 Mbps 1500

```

```

Interface Vsan Admin Status Oper Oper
 Mode Trunk Mode Mode Speed
 (Gbps)

port-channel 1 1 on trunking TE 4
port-channel 2 1 on trunking TE 4

```

```

Interface Vsan Admin Admin Status Oper Profile Port-channel
 Mode Mode Trunk Mode Mode
 Mode

fcip10 1 auto on notConnected -- 10 --

```

### Example 10-7 Displays Interface Counters

```

switch# show interface counters
fc3/1
 5 minutes input rate 24 bits/sec, 3 bytes/sec, 0 frames/sec
 5 minutes output rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
 3502 frames input, 268400 bytes
 0 discards, 0 CRC, 0 unknown class
 0 too long, 0 too short
 3505 frames output, 198888 bytes
 0 discards

```

```

1 input OLS, 1 LRR, 1 NOS, 0 loop inits
2 output OLS, 1 LRR, 1 NOS, 0 loop inits
1 link failures, 1 sync losses, 1 signal losses
.
.
.
fc9/8
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
0 frames input, 0 bytes
0 class-2 frames, 0 bytes
0 class-3 frames, 0 bytes
0 class-f frames, 0 bytes
0 discards, 0 CRC, 0 unknown class
0 too long, 0 too short
0 frames output, 0 bytes
0 class-2 frames, 0 bytes
0 class-3 frames, 0 bytes
0 class-f frames, 0 bytes
0 discards
0 input OLS, 0 LRR, 0 NOS, 0 loop inits
0 output OLS, 0 LRR, 0 NOS, 0 loop inits
0 link failures, 0 sync losses, 0 signal losses
16 receive B2B credit remaining
3 transmit B2B credit remaining.
. . .
sup-fc0
114000 packets input, 11585632 bytes
0 multicast frames, 0 compressed
0 input errors, 0 frame, 0 overrun 0 fifo
113997 packets output, 10969672 bytes, 0 underruns
0 output errors, 0 collisions, 0 fifo
0 carrier errors

mgmt0
31557 packets input, 2230860 bytes
0 multicast frames, 0 compressed
0 input errors, 0 frame, 0 overrun 0 fifo
26618 packets output, 16824342 bytes, 0 underruns
0 output errors, 0 collisions, 7 fifo
0 carrier errors

vsan1
0 packets input, 0 bytes, 0 errors, 0 multicast
0 packets output, 0 bytes, 0 errors, 0 dropped
.
.
.
port-channel 1
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
0 frames input, 0 bytes
0 class-2 frames, 0 bytes
0 class-3 frames, 0 bytes
0 class-f frames, 0 bytes
0 discards, 0 CRC, 0 unknown class
0 too long, 0 too short
0 frames output, 0 bytes
0 class-2 frames, 0 bytes
0 class-3 frames, 0 bytes
0 class-f frames, 0 bytes
0 discards
0 input OLS, 0 LRR, 0 NOS, 0 loop inits
0 output OLS, 0 LRR, 0 NOS, 0 loop inits

```

0 link failures, 0 sync losses, 0 signal losses



**Note** Interfaces 9/8 and 9/9 are not trunking ports and display class 2, 3, and F information as well.

### Example 10-8 Displays Interface Counters in Brief Format

```
switch# show interface counters brief
```

```

Interface Input (rate is 5 min avg) Output (rate is 5 min avg)

Rate Total Rate Total
Mbits/s Frames Mbits/s Frames

fc3/1 0 3871 0 3874
fc3/2 0 3902 0 4232
fc3/3 0 3901 0 4138
fc3/4 0 3895 0 3894
fc3/5 0 3890 0 3897
fc9/8 0 0 0 0
fc9/9 0 5 0 4
fc9/10 0 4186 0 4182
fc9/11 0 4331 0 4315

```

```

Interface Input (rate is 5 min avg) Output (rate is 5 min avg)

Rate Total Rate Total
Mbits/s Frames Mbits/s Frames

port-channel 1 0 0 0 0
port-channel 2 0 3946 0 3946

```

### Example 10-9 Displays BB\_credit Information

```
switch# show interface bccredit
fc2/1 is down (Fcot not present)
...
fc2/17 is trunking
 Transmit B2B Credit is 255
 Receive B2B Credit is 12
 Receive B2B Credit performance buffers is 375
 12 receive B2B credit remaining
 255 transmit B2B credit remaining
fc2/18 is down (Fcot not present)
fc2/19 is down (Fcot not present)
fc2/20 is down (Fcot not present)
fc2/21 is down (Link failure or not-connected)
...
fc2/31 is up
 Transmit B2B Credit is 0
 Receive B2B Credit is 12
 Receive B2B Credit performance buffers is 48
 12 receive B2B credit remaining
 0 transmit B2B credit remaining
fc2/32 is down (Link failure or not-connected)
```



**Example 10-10 Displays BB\_credit Information for a Specified Fibre Channel Interface**

```
switch# show interface fc2/31 bbcredit
fc2/31 is up
 Transmit B2B Credit is 0
 Receive B2B Credit is 12
 Receive B2B Credit performance buffers is 48
 12 receive B2B credit remaining
 0 transmit B2B credit remaining
```

**Note**

The **show interface interface-type slot/port transceiver** command can only be issued on a switch in the Cisco MDS 9100 Series if the FCOT is present (see [Example 10-11](#) to [Example 10-14](#)).

**Example 10-11 Displays Transceiver Information**

```
switch# show interface transceiver
fc1/1 fcot is present
 name is CISCO-AGILENT
 part number is QFBR-5796L
 revision is
 serial number is A00162193
 fc-transmitter type is short wave laser
 cisco extended id is unknown (0x0)

...
fc1/9 fcot is present
 name is FINISAR CORP.
 part number is FTRJ-1319-7D-CSC
 revision is
 serial number is H11A6ER
 fc-transmitter type is long wave laser cost reduced
 cisco extended id is unknown (0x0)

...
```

[Example 10-12](#) displays diagnostics information for a specified Fibre Channel interface.

**Example 10-12 Displays Transceiver Information for a Specified Fibre Channel Interface**

```
switch# show interface fc1/9 transceiver
fc1/9 fcot is present
 name is CISCO-FINISAR
 part number is FTRJ8519P1BNL-C3
 revision is A
 serial number is FNS0743Z00D
 fc-transmitter type is short wave laser
 cisco extended id is unknown (0x0)

SFP Diagnostics Information
 Temperature : 43.70 Celsius
 Voltage : 3.28 Volt
 Current : 5.23 mA
 Optical Tx Power : 8.32 dBm
 Optical Rx Power : -4.95 dBm
Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning
```

[Example 10-13](#) displays the real-time diagnostics information for transceivers, including the temperature, voltage, current, tx/rx-power. It also shows the hi/lo alarm and warning thresholds.

**Example 10-13 Displays Transceiver Information for a Specified Fibre Channel Interface**

```

switch# show interface fc1/9 transceiver detail
fc1/9 fcot is present
 name is CISCO-FINISAR
 part number is FTRJ8519P1BNL-C3
 revision is A
 serial number is FNS0743Z00D
 fc-transmitter type is short wave laser
 cisco extended id is unknown (0x0)

 SFP Detail Diagnostics Information

 Alarms Warnings
 High Low High Low

Temperature 43.70 C 109.00 C -29.00 C 103.00 C -13.00 C
Voltage 3.28 V 3.90 V 2.70 V 3.70 V 2.90 V
Current 5.23 mA 15.00 mA 1.00 mA 12.00 mA 2.00 mA
Tx Power 8.32 dBm 8.56 dBm 8.20 dBm 8.56 dBm 8.22 dbm
Rx Power -4.95 dBm 1.00 dBm -20.00 dBm -1.00 dBm -18.02 dBm

Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning

```

[Example 10-14](#) displays the calibrations used for computing the diagnostics information. If the SFP is internally calibrated, the output provides this information. Otherwise, it provides various (standard) calibration constants.

**Tip**

This command can only be issued on a switch in the Cisco MDS 9100 Series if the FCOT is present.

**Example 10-14 Displays Transceiver Information for a Specified Fibre Channel Interface**

```

switch# show interface fc1/9 transceiver calibrations
fc1/9 fcot is present
 name is CISCO-FINISAR
 part number is FTRJ8519P1BNL-C3
 revision is A
 serial number is FNS0743Z00D
 fc-transmitter type is short wave laser
 cisco extended id is unknown (0x0)

 SFP Calibrations Information

 Slope Offset Rx4/Rx3/Rx2/Rx1/Rx0

Temperature 256 0
Voltage 256 0
Current 256 0
Tx Power 143 65508
Rx Power 0.0000/ 0.0000/ 0.0000/ 0.2057/ -7.6935

```

[Example 10-15](#) displays the running configuration for a specified interface.

**Example 10-15 Displays the Running Configuration for a Specified Interface**

```

switch# show running-config interface fc1/1
interface fc1/1
switchport mode FL
no shutdown

```

## Displaying TL Port Information

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric since they only communicate with devices on the same physical loop.

The legacy devices are used in Fibre Channel networks and devices outside the loop may need to communicate with them. The communication functionality is provided through TL ports.

Use the **switchport mode** command to configure a TL port (see the [“Configuring Interface Modes” section on page 10-9](#)).

The **show tlport** command displays the TL port interface configurations. This command provides a list of all TL ports configured on a box and shows the associated VSAN, the FC ID for the port (only domain and area are valid), and the current operational state of the TL port (up or initializing). See Examples 10-16 to 10-19.

### Example 10-16 Displays the TL Ports in All VSANs

```
switch# show tlport list

Interface Vsan FC-ID State

fc1/16 1 0x420000 Init
fc2/26 1 0x150000 Up
```

TL ports allow a private device (devices that physically reside on the loop) to see a fabric device and vice-versa by proxying fabric devices on the loop. Fabric devices are proxied by allocating each fabric device an ALPA on this loop.

In addition to these proxied devices, other virtual devices (local or remote domain controller addresses) are also allocated ALPAs on the loop. A switch reserves the ALPA for its own communication with private devices, and the switch acts as a SCSI Initiator.

The first column in the output of the **show tlport interface** command is the ALPA identity of the device on the loop. The second lists the port WWNs, the third lists the node WWNs for each device, the fourth identifies the device as a SCSI initiator or target, and the last column is the real FC ID of the device.

### Example 10-17 Displays the Detailed Information for a Specific TL Port

```
switch# show tlport interface fc1/16 all
fc1/16 is up, vsan 1, FCID 0x420000

alpa pWWN nWWN SCSI Type Device FC-ID

0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Proxied 0xfffc42
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target Private 0x420073
0xef 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Switch 0x0000ef
```

### Example 10-18 Displays TL Port Information for Private Devices

```
switch# show tlport int fc1/16 pri
fc1/16 is up, vsan 1, FCID 0x420000

alpa pWWN nWWN SCSI Type FC-ID

0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target 0x420073
```

```
0x74 22:00:00:20:37:38:d3:de 20:00:00:20:37:38:d3:de Target 0x420074
```

**Example 10-19 Displays TL Port Information for Proxied Devices**

```
switch# show tlport int fc1/16 prox
fc1/16 is up, vsan 1, FCID 0x420000
```

```

alpa pWWN nWWN SCSI Type FC-ID

0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator 0xffffc42
0x02 21:00:00:e0:8b:01:95:e7 20:00:00:e0:8b:01:95:e7 Initiator 0x420100
```

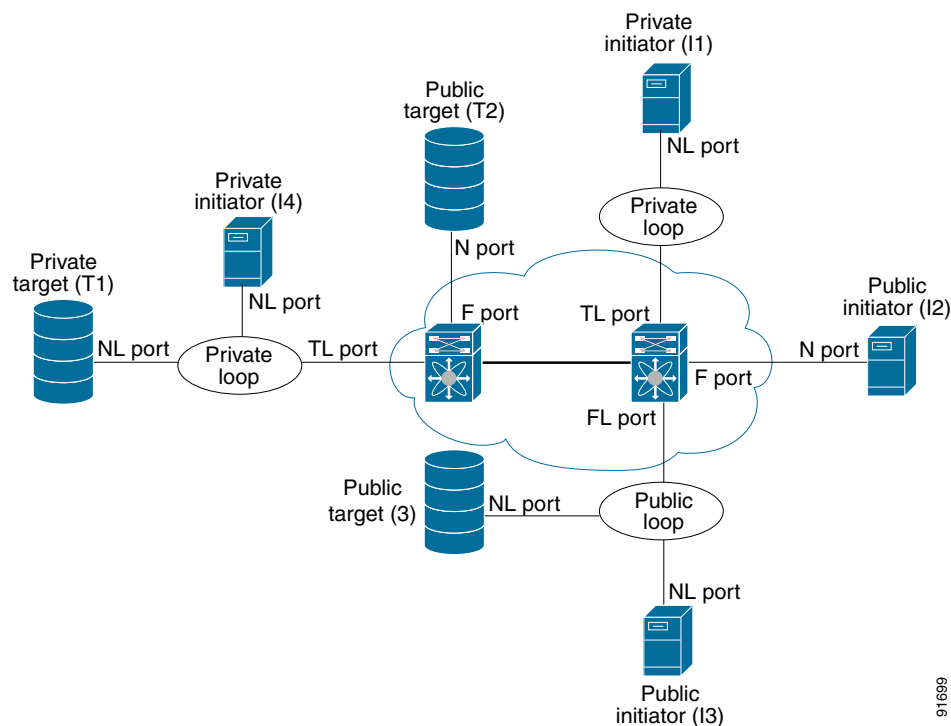
## TL Port Translation Guidelines

Table 10-7 lists the TL port translations supported in Cisco MDS 9000 Family switches:

**Table 10-7 Supported TL Port Translation**

| Translation from           | Translation to          | Example (see Figure 10-3)   |
|----------------------------|-------------------------|-----------------------------|
| Private initiator          | Private target          | From I1 to T1 or vice versa |
| Private initiator          | Public target — N port  | From I1 to T2 or vice versa |
| Private initiator          | Public target — NL port | From I4 to T3 or vice versa |
| Public initiator — N port  | Private target          | From I2 to T1 or vice versa |
| Public initiator — NL port | Private target          | From I3 to T1 or vice versa |

**Figure 10-3 TL Port Translation Support Examples**



Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxied to a private loop.
- To be proxied to the private loop, fabric devices must be in the same zone as private loop devices.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.
- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.





# Configuring Trunking

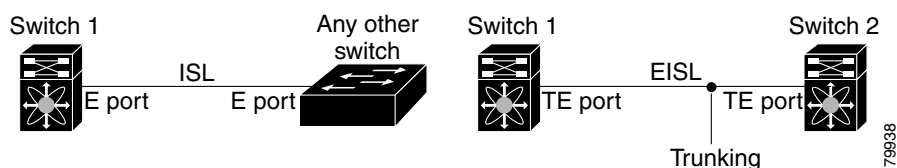
This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

- [About Trunking, page 11-1](#)
- [About Trunking Protocol, page 11-2](#)
- [Configuring Trunk Modes, page 11-3](#)
- [Configuring Trunk-Allowed VSAN List, page 11-4](#)
- [Trunking Configuration Guidelines, page 11-6](#)
- [Displaying Trunking Information, page 11-7](#)
- [Default Settings, page 11-8](#)

## About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using Extended ISL (EISL) frame format (see [Figure 11-1](#)).

**Figure 11-1 Trunking**



The trunking feature includes the following restrictions:

- Trunking configurations are only applicable to E ports. If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port (see the [“Configuring Trunk Modes”](#) section on page 11-3).
- The trunk-allowed VSANs configured for TE ports are used by the trunking protocol to determine the allowed-active VSANs in which frames can be received or transmitted (see the [“Configuring Trunk-Allowed VSAN List”](#) section on page 11-4).
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port (see the [“About Trunking Protocol”](#) section on page 11-2).

# About Trunking Protocol

The trunking protocol is important for E-port and TE-port operations. It supports the following:

- Dynamic negotiation of operational trunk mode (see the [“Configuring Trunk Modes” section on page 11-3](#)).
- Selection of a common set of trunk-allowed VSANs (see the [“Configuring Trunk-Allowed VSAN List” section on page 11-4](#)).
- Detection of a VSAN mismatch across an ISL (see the [“Trunking Configuration Guidelines” section on page 11-6](#)).

By default, the trunking protocol is enabled. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations will not be affected—the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

**Tip**

To avoid inconsistent configurations, shut all E ports before enabling or disabling the trunking protocol.

To enable or disable the trunking protocol, follow these steps:

|        | Command                                                            | Purpose                              |
|--------|--------------------------------------------------------------------|--------------------------------------|
| Step 1 | switch# <b>config t</b>                                            | Enters configuration mode.           |
| Step 2 | switch(config)# <b>no trunk protocol enable</b><br>switch(config)# | Disables the trunking protocol.      |
|        | switch(config)# <b>trunk protocol enable</b><br>switch(config)#    | Enables trunking protocol (default). |



# Configuring Trunk Modes

By default, the trunk mode is enabled in all Fibre Channel interfaces. However, the trunk mode configuration takes effect only in E-port mode. You can configure the trunk mode as **on** (enabled), **off** (disabled), or **auto** (automatic). The default trunk mode is **on**. The trunk mode configuration at the two ends of an ISL, between two switches, determine the resulting trunking state of the link and the port modes at both ends (see [Table 11-1](#)).

**Table 11-1 Trunk Mode Status Between Switches**

| Your Trunk Mode Configuration |                  | Resulting State and Port Mode |           |
|-------------------------------|------------------|-------------------------------|-----------|
| Switch 1                      | Switch 2         | Trunking State                | Port Mode |
| On                            | Auto or on       | Trunking (EISL)               | TE port   |
| Off                           | Auto, on, or off | No trunking (ISL)             | E port    |
| Auto                          | Auto             | No trunking (ISL).            | E port    |



## Note

When connected to a third-party switch, the trunk mode configuration has no effect—the ISL is always in a trunking disabled state.

To configure the trunk mode, follow these steps:

|        | Command                                                                    | Purpose                                                                                                                     |
|--------|----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                    | Enters configuration mode.                                                                                                  |
| Step 2 | switch(config)# <b>interface fc1/1</b><br>switch(config-if)#               | Configures the specified interface.                                                                                         |
| Step 3 | switch(config-if)# <b>switchport trunk mode on</b><br>switch(config-if)#   | Enables the trunk mode for the specified interface.                                                                         |
|        | switch(config-if)# <b>switchport trunk mode off</b><br>switch(config-if)#  | Disables the trunk mode for the specified interface.                                                                        |
|        | switch(config-if)# <b>switchport trunk mode auto</b><br>switch(config-if)# | Configures the trunk mode for the specified interface. The <b>auto</b> option provides automatic sensing for the interface. |

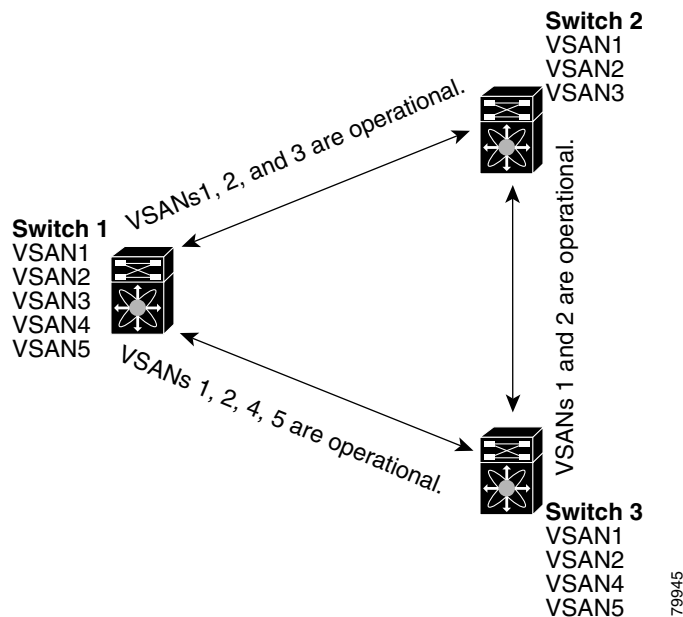
## Configuring Trunk-Allowed VSAN List

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

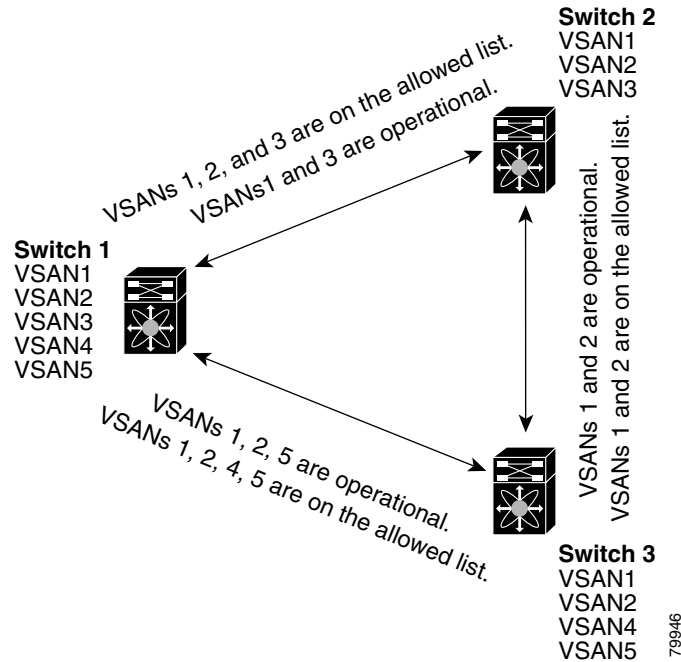
In [Figure 11-2](#), switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in [Figure 11-2](#).

**Figure 11-2 Default Allowed -Active VSAN Configuration**



You can configure a select set of VSANs (from the allowed-active list) to control access to those VSANs in a trunking ISL. Using Figure 11-2 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 11-3).

**Figure 11-3 Operational and Allowed VSAN Configuration**



In Figure 11-3, the operational allowed list of VSANs between switches is as follows:

- Switch 1 and switch 2 include VSAN 1 and VSAN 3.
- Switch 2 and switch 3 include VSAN 1 and VSAN 2.
- Switch 3 and switch 1 include VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

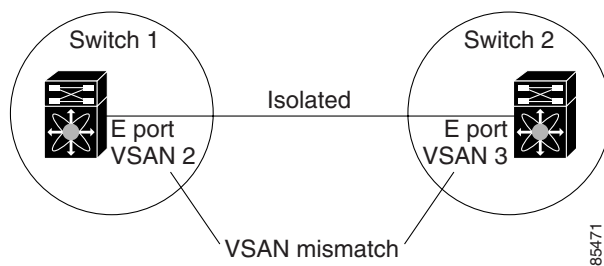
To configure an allowed-active list of VSANs for an interface, follow these steps:

|               | Command                                                                                            | Purpose                                                 |
|---------------|----------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| <b>Step 1</b> | switch# <code>config t</code>                                                                      | Enters configuration mode.                              |
| <b>Step 2</b> | switch(config)# <code>interface fc1/1</code><br>switch(config-if)#                                 | Configures the specified interface.                     |
| <b>Step 3</b> | switch(config-if)# <code>switchport trunk allowed vsan 2-4</code>                                  | Changes the allowed list for the specified VSANs.       |
|               | switch(config-if)# <code>switchport trunk allowed vsan add 5</code><br>updated trunking membership | Expands the specified VSAN (5) to the new allowed list. |
|               | switch(config-if)# <code>no switchport trunk allowed vsan 2-4</code>                               | Deletes VSANs 2, 3, and 4.                              |
|               | switch(config-if)# <code>no switchport trunk allowed vsan add 5</code>                             | Deletes the expanded allowed list.                      |

# Trunking Configuration Guidelines

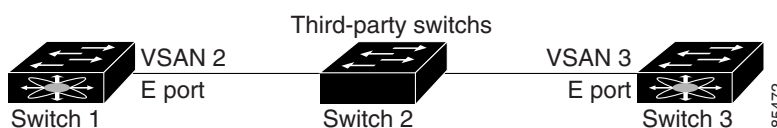
If you misconfigure VSAN configurations across E ports, you could face consequences such as merging the traffic in two VSANs (thus causing both VSANs to mismatch). The trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid VSANs merging (see [Figure 11-4](#)).

**Figure 11-4 VSAN Mismatch**



In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved. The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see [Figure 11-5](#)).

**Figure 11-5 Third-Party Switch VSAN Mismatch**



VSANs 2 and 3 get effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies (refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*).

# Displaying Trunking Information

The **show interface** command is invoked from the EXEC mode and displays trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch. See Examples 11-1 to 11-3.

## Example 11-1 Displays a Trunked Fiber Channel Interface

```
switch# show interface fc1/13
fc1/13 is trunking
 Hardware is Fibre Channel
 Port WWN is 20:0d:00:05:30:00:58:1e
 Peer port WWN is 20:0d:00:05:30:00:59:1e
 Admin port mode is auto, trunk mode is on
 Port mode is TE
 Port vsan is 1
 Speed is 2 Gbps
 Receive B2B Credit is 255
 Beacon is turned off
 Trunk vsans (admin allowed and active) (1)
 Trunk vsans (up) (1)
 Trunk vsans (isolated) ()
 Trunk vsans (initializing) ()
 5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 233996 frames input, 14154208 bytes, 0 discards
 0 CRC, 0 unknown class
 0 too long, 0 too short
 236 frames output, 13818044 bytes, 0 discards
 11 input OLS, 12 LRR, 10 NOS, 28 loop inits
 34 output OLS, 19 LRR, 17 NOS, 12 loop inits
```

## Example 11-2 Displays Trunking Protocol

```
switch# show trunk protocol
Trunk protocol is enabled
```

## Example 11-3 Displays Per VSAN Information on Trunk Ports

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/7 is trunking
 Vsan 1000 is down (Isolation due to vsan not configured on peer)
...
fc3/10 is trunking
 Vsan 1 is up, FCID is 0x760001
 Vsan 2 is up, FCID is 0x6f0001

fc3/11 is trunking
 Belongs to port-channel 6
 Vsan 1 is up, FCID is 0xef0000
 Vsan 2 is up, FCID is 0xef0000
...
port-channel 6 is trunking
 Vsan 1 is up, FCID is 0xef0000
 Vsan 2 is up, FCID is 0xef0000
```

# Default Settings

Table 11-2 lists the default settings for trunking parameters.

**Table 11-2 Default Trunk Configuration Parameters**

| Parameters             | Default                         |
|------------------------|---------------------------------|
| Switch port trunk mode | On                              |
| Allowed VSAN list      | 1 to 4093 user-defined VSAN IDs |
| Trunking protocol      | Enabled                         |



## Configuring PortChannels

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link. Specifically, a PortChannel has the following functionality:

- Provides a point-to-point connection over an ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.



### Note

See the [“Fail-over Scenarios for PortChannels and FSPF Links”](#) section on [page 19-3](#) for fail-over scenarios.

Cisco MDS 9000 Family of switches support 128 PortChannels with 16 interfaces per PortChannel.

This chapter discusses the PortChannel feature provided in the switch. This chapter includes the following sections:

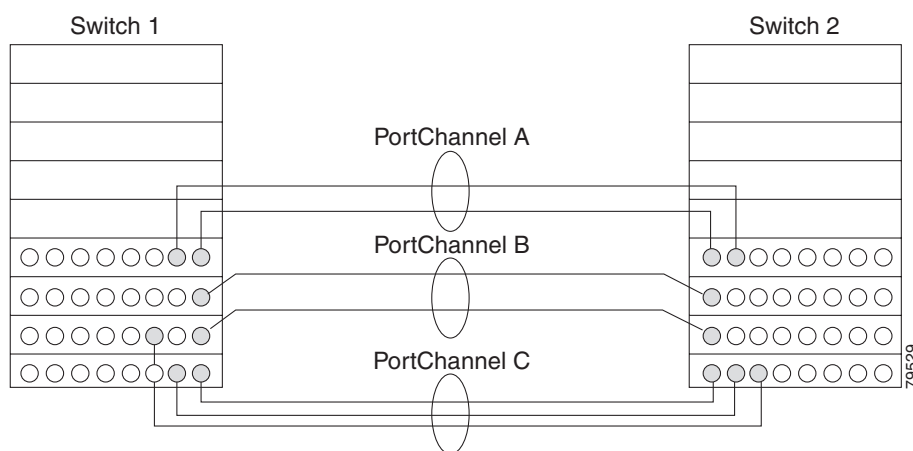
- [PortChannel Examples, page 12-2](#)
- [About PortChanneling and Trunking, page 12-3](#)
- [About Load Balancing, page 12-4](#)
- [Creating a PortChannel, page 12-5](#)
- [Deleting a PortChannel, page 12-6](#)
- [Adding Interfaces to a PortChannel, page 12-6](#)
- [Deleting Interfaces from a PortChannel, page 12-8](#)
- [Considerations for PortChannel Configurations, page 12-9](#)
- [Viewing PortChannel Information, page 12-10](#)
- [Default Settings, page 12-12](#)

# PortChannel Examples

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. [Figure 12-1](#) illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

**Figure 12-1 PortChannel Flexibility**



## Configuring 32-port Switching Modules and Host-Optimized Ports

The 32-port 1/2-Gbps switching module contains 8 port groups of 4 ports each. When configuring these modules or the host-optimized ports in the Cisco 9100 Series, the following guidelines apply:

- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to same PortChannel rules as 32-port switching modules—only the first port of each group of 4 ports is included in a PortChannel.
  - You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain in the shutdown state.
  - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port as a PortChannel. The other three ports continue to remain in a no shutdown state.



### Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.



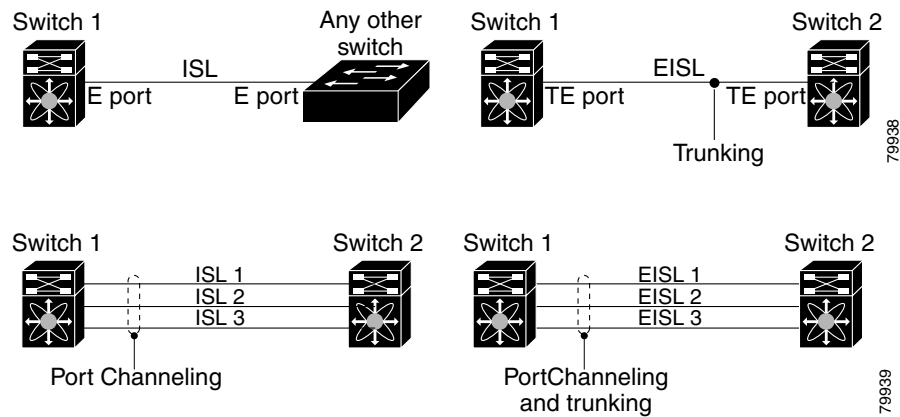
# About PortChanneling and Trunking

PortChanneling enables several links to be combined into one aggregated link.

Trunking enables an ISL to carry (trunk) multiple VSANs. Trunking can only be configured on a TE port. A TE port is specific to switches in the Cisco MDS 9000 Family. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (see [Figure 12-2](#)).

See [Chapter 11, “Configuring Trunking”](#) for information on trunked interfaces.

**Figure 12-2 PortChanneling and Trunking**



PortChanneling and trunking are used separately across an ISL:

- PortChanneling—Interfaces can be channeled between E ports over multiple ISLs or between TE ports over multiple EISLs.
- Trunking—Trunking, which permits carrying VSAN IDs between switches, can be done only between TE ports over EISLs.

See [Chapter 9, “Configuring and Managing VSANs.”](#)

Both PortChanneling and trunking can be used between TE ports over EISLs.

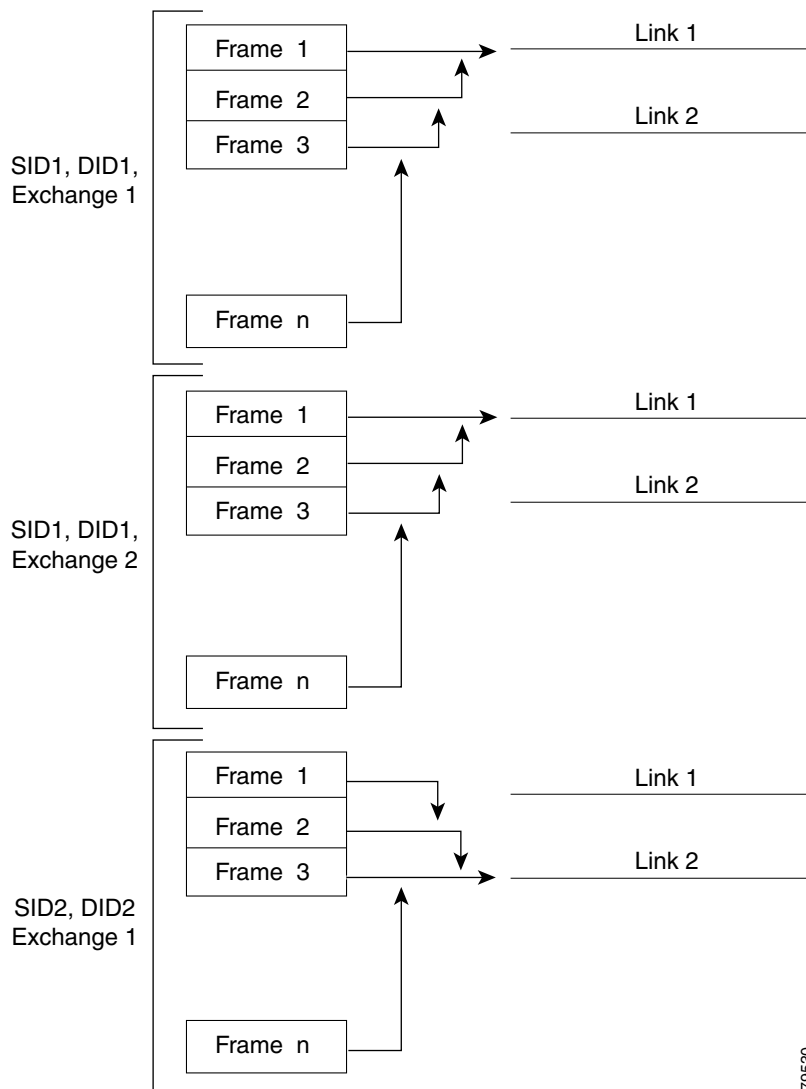
# About Load Balancing

Two mechanisms support the load balancing functionality:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange picks a link and subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This provides more granular load balancing while preserving the order of frames for each exchange.

Figure 12-3 illustrates how source ID 1 (SID1) and destination ID1-based(DID1) load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

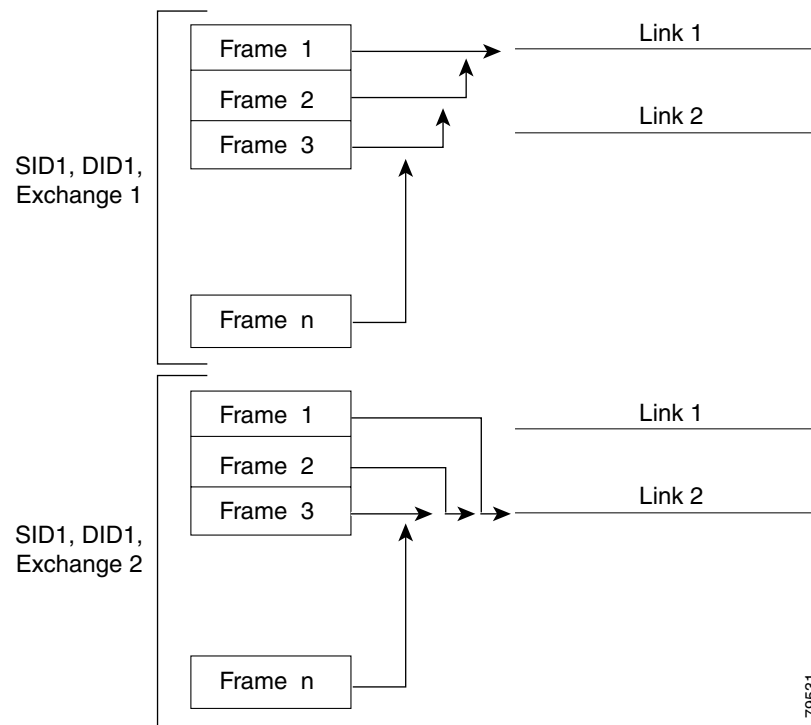
**Figure 12-3 SID1 and DID1Based Load Balancing**



79530

Figure 12-4 illustrates how exchange based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

**Figure 12-4 SID1, DID1, and Exchange Based Load Balancing**



For more information on configuring load balancing and in-order delivery features, see the “[VSAN Attributes](#)” section on page 9-6.

## Creating a PortChannel

You can create PortChannels using the **interface port-channel** command. PortChannels are created with default values. You can change the default configuration just like any other physical interface.

To create a PortChannel, follow these steps:

|        | Command                                                               | Purpose                                   |
|--------|-----------------------------------------------------------------------|-------------------------------------------|
| Step 1 | switch# <b>conf t</b>                                                 | Enters configuration mode.                |
| Step 2 | switch(config)# <b>interface port-channel 1</b><br>switch(config-if)# | Configures the specified PortChannel (1). |



### Note

All interfaces added to PortChannels are administratively shut down, and the PortChannel remains administratively up.

## Deleting a PortChannel

To delete the PortChannel, you must explicitly issue the **no interface port-channel** command. When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. To avoid inconsistent states across switches, and to maintain consistency across switches, the ports shut down. They continue to use the configured values of the physical port.

To delete a PortChannel, follow these steps:

|        | Command                                                                                                                                                                                                         | Purpose                                                                                                                       |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                         | Enters configuration mode.                                                                                                    |
| Step 2 | switch(config)# <b>no interface port-channel 1</b><br>port-channel 1 deleted and all its members disabled<br>please do the same operation on the switch at the other end of the port-channel<br>switch(config)# | Deletes the specified PortChannel (1), its associated interface mappings, and the hardware associations for this PortChannel. |

## Adding Interfaces to a PortChannel

You can add a physical interface (or a range of interfaces) to a nonexistent or an existing PortChannel and the PortChannel is automatically created. If the PortChannel does not exist, it is created. The compatible parameters on the configuration are mapped to the PortChannel.

To add a port (or a range of ports) to a PortChannel, follow these steps:

|        | Command                                                                                                                                                                                                                                                                       | Purpose                                                                                                                                                                                                                                                            |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                                                                                       | Enters configuration mode.                                                                                                                                                                                                                                         |
| Step 2 | switch(config)# <b>interface fc1/15</b><br>switch(config-if)#                                                                                                                                                                                                                 | Configures the specified port interface (fc1/15).                                                                                                                                                                                                                  |
|        | switch(config)# <b>interface fc1/1 - 5</b><br>switch(config-if)#                                                                                                                                                                                                              | Configures the specified range of interfaces. In this example, interfaces from 1/1 to 1/5                                                                                                                                                                          |
| Step 3 | switch(config-if)# <b>channel-group 15</b><br>fc1/15 added to port-channel 15 and disabled<br>please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up<br>switch(config-if)#                      | Adds physical Fibre Channel port 1/15 to channel group 15. If channel group 15 does not exist, it is created. The port is shut down.                                                                                                                               |
|        | switch(config-if)# <b>channel-group 2</b><br>fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 added to port-channel 2 and disabled<br>please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up<br>switch(config-if)# | Adds physical interfaces 1/1, 1/2, 1/3, 1/4, and 1/5 to channel group 2. If channel group 2 does not exist, it is created.<br><br>If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces. |

## Forcing an Interface Addition

You can specify a **force** option to force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel and the port is shut down.



### Note

When PortChannels are created automatically, the **force** option cannot be used.

To force the addition of a port to a PortChannel, follow these steps:

|        | Command                                                                                                                                                                                                                                                           | Purpose                                                                                              |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                                                                           | Enters configuration mode.                                                                           |
| Step 2 | switch(config)# <b>interface fc1/1</b><br>switch(config-if)#                                                                                                                                                                                                      | Configures the specified port interface (fc1/1).                                                     |
| Step 3 | switch(config-if)# <b>channel-group 1 force</b><br>fc1/1 added to port-channel 1 and disabled<br>please do the same operation on the switch at<br>the other end of the port-channel, then do<br>"no shutdown" at both ends to bring them up<br>switch(config-if)# | Forces a physical Fibre Channel port 1/1<br>addition to channel group 1. The E port is shut<br>down. |

## Compatibility Check

A compatibility check ensures that the same configuration values are used in all physical ports in the channel. For example, to enable trunk mode, all operational ports in the configuration must be configured in the trunk mode or in the nontrunking mode. Otherwise, they cannot become part of a PortChannel. A port cannot be operational if it is incompatible with the PortChannel. If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces.

## Suspended State

An interface enters the suspended state if its operational values are incompatible with the PortChannel. A compatibility check on operational parameters is done when one of the following events occurs:

- A port becomes operational in a PortChannel.
- An operational parameter changes for a port in a PortChannel.

The software performs a compatibility check on the operational parameters and places the interface in an operational or suspended state based on the result (see the [“Reason Codes”](#) section on page 10-6).

# Deleting Interfaces from a PortChannel

To delete a physical interface (or a range of physical interfaces), you must explicitly issue the **no channel-group** command at the physical interface level. When a physical interface is deleted from the PortChannel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the PortChannel status is changed to a down state. Deleting an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.



## Note

When an interface is deleted, it is shut down but the physical configuration is retained. The inherited PortChannel configuration information is not deleted.

To delete a physical interface (or a range of physical interfaces), follow these steps:

|        | Command                                                                                                                                                                                                                                | Purpose                                                              |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| Step 1 | switch(config)# <b>interface</b> fc1/1<br>switch(config-if)#                                                                                                                                                                           | Enters the selected physical interface level.                        |
|        | switch(config)# <b>interface</b> fc1/1 - 5<br>switch(config-if)#                                                                                                                                                                       | Enters the selected range of physical interfaces.                    |
| Step 2 | switch(config-if)# <b>no channel-group</b> 2<br>fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 removed<br>from port-channel 2 and disabled. Please<br>do the same operation on the switch at<br>the other end of the port-channel<br>switch(config-if)# | Deletes the physical Fibre Channel interfaces in<br>channel group 2. |

# Quiescing a PortChannel ISL

Generally a **shutdown** command issued on an interface through which traffic is flowing disables the interface with possible frame drop. You can avoid this frame drop, by using the **quiesce** command to gracefully shutdown an interface without dropping any frames. This command can only be issued on an ISL within a PortChannel—at both ends of the link. This addition prevents frame loss for a planned link shutdown or removal.



## Note

If the in-order delivery feature is enabled, some frames may still be dropped. By default, the in-order delivery feature is disabled in all switches in the Cisco MDS 9000 Family.

The following conditions return an error:

- The interface is not part of port-channel
- The interface is not up
- The interface is the last operational interface in the PortChannel

You can negate a previously-issued **quiesce** command by issuing the **quiesce no interface** before the interface is shut down.

To gracefully shutdown an ISL in a PortChannel, follow these steps.

**Step 1** Log in to Switch A and issue the **quiesce** command on the ISL.

```
switchA# quiesce interface fc 2/1
```

```
WARNING: this command will stop forwarding frames to the specified interfaces. It is
intended to be used to gracefully shutdown interfaces in a port-channel. The procedure is:
1. quiesce the interfaces on both switches.
2. shutdown the interfaces administratively.
Do you want to continue? (y/n) [n] y
```

**Step 2** Log in to Switch N and issue the **quiesce** command on the ISL.

```
switchN# quiesce interface fc 2/9
WARNING: this command will stop forwarding frames to the specified interfaces. It is
intended to be used to gracefully shutdown interfaces in a port-channel. The procedure is:
1. quiesce the interfaces on both switches.
2. shutdown the interfaces administratively.
Do you want to continue? (y/n) [n] y
```

**Step 3** Ensure that the command has completed its sequence and issue the shutdown command on one or both switches.

```
switchN(config-if)# shutdown
```

---

## Considerations for PortChannel Configurations

Before configuring a PortChannel, consider the following guidelines

- Configure the PortChannel across switching modules to prevent redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to two switches. PortChannels require point-to-point connections.

## Error Detection

If you misconfigure PortChannels, you may receive the `Error disabled - Possible port channel misconfiguration` message. If you receive this message, the PortChannel's physical links are disabled since an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side.
- Links in a PortChannel must not be changed after the PortChannel is configured.

If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and re-enable the links. Issue the **show interface** command for that interface to verify that the PortChannel is functioning as required.

If all three conditions are not met, the faulty link is disabled.

# Viewing PortChannel Information

You can view specific information about existing PortChannels at any time from EXEC mode. The following **show** commands provide further details on existing PortChannels. You can force all screen output to go to a printer or save it to a file.

The **show port-channel summary** command displays a summary of PortChannels within the switch. A one-line summary of each PortChannel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational interface (FOP), which is the primary operational interface selected in the PortChannel. See Examples 12-1 to 12-6.

## Example 12-1 PortChannel Summary

```
switch# show port-channel summary
```

| Interface      | Total Ports | Oper Ports | First Oper Port |
|----------------|-------------|------------|-----------------|
| port-channel 1 | 2           | 2          | fc2/3           |
| port-channel 2 | 2           | 2          | fc2/5           |
| port-channel 3 | 2           | 2          | fc2/10          |
| .              |             |            |                 |
| .              |             |            |                 |
| .              |             |            |                 |

## Example 12-2 PortChannel Compatibility

```
switch# show port-channel compatibility-parameters
physical port layer fibre channel or ethernet
 port mode E/AUTO only
 trunk mode
 speed
 port VSAN
 port allowed VSAN list
```

## Example 12-3 PortChannel Database

```
switch# show port-channel database
port-channel 1
 Administrative channel mode is on
 Operational channel mode is on
 Last membership update succeeded
 First operational port is fc2/3
 2 ports in total, 2 ports up
 Ports: fc2/3 [up]
 fc2/4 [up]
port-channel 2
 Administrative channel mode is on
 Operational channel mode is on
 Last membership update succeeded
 First operational port is fc2/5
 2 ports in total, 2 ports up
 Ports: fc2/5 [up]
 fc2/6 [up]
.
.
.
```



The **show port-channel consistency** command has two options—without detail and **detail**.

#### Example 12-4 Command Without Details

```
switch# show port-channel consistency
sup database:
=====
totally 7 port-channels
port-channel 1:
 2 ports, first operational port is fc2/3
 fc2/3 [up]
 fc2/4 [up]
port-channel 2:
 2 ports, first operational port is fc2/5
 fc2/5 [up]
 fc2/6 [up]
.
.
.
```

#### Example 12-5 Command With Details

```
switch# show port-channel consistency detail
Authoritative port-channel database:
=====
totally 7 port-channels
port-channel 1:
 2 ports, first operational port is fc2/3
 fc2/3 [up]
 fc2/4 [up]
port-channel 2:
 2 ports, first operational port is fc2/5
 fc2/5 [up]
 fc2/6 [up]
.
.
.
=====
database 1: from module 5
=====
totally 7 port-channels
port-channel 1:
 2 ports, first operational port is fc2/3
 fc2/3 [up]
 fc2/4 [up]
port-channel 2:
 2 ports, first operational port is fc2/5
 fc2/5 [up]
 fc2/6 [up]
.
.
.
=====
database 3: from module 2
=====
totally 7 port-channels
port-channel 1:
 2 ports, first operational port is fc2/3
 fc2/3 [up]
 fc2/4 [up]
port-channel 2:
 2 ports, first operational port is fc2/5
```

```
fc2/5 [up]
fc2/6 [up]
.
.
.
```

The **show port-channel usage** command displays details of the used and unused PortChannel numbers.

**Example 12-6 PortChannel Usage**

```
switch# show port-channel usage
Totally 7 port-channel numbers used
=====
Used : 1-7
Unused: 8-128
```

# Default Settings

[Table 12-1](#) lists the default settings for PortChannels.

**Table 12-1 Default PortChannel Parameters**

| Parameters         | Default                     |
|--------------------|-----------------------------|
| PortChannels       | FSPF is enabled by default. |
| Create PortChannel | Administratively up.        |
| Default mode       | Auto.                       |
| Quiesce            | Disabled.                   |



## Configuring and Managing Zones

---

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field. This chapter defines various zoning concepts and provides details on zone set and management features in the switch and includes the following sections:

- [Zoning Features, page 13-2](#)
- [Zoning Example, page 13-3](#)
- [Configuring a Zone, page 13-4](#)
- [Configuring Aliases, page 13-5](#)
- [Zone Enforcement, page 13-9](#)
- [Zone Sets, page 13-6](#)
- [The Default Zone, page 13-9](#)
- [Recovering from Link Isolation, page 13-10](#)
- [Distributing Zone Sets, page 13-11](#)
- [Copying Zone Sets, page 13-11](#)
- [Clearing the Zone Database, page 13-12](#)
- [LUN Zoning, page 13-12](#)
- [Read-Only Zoning, page 13-13](#)
- [Viewing Zone Information, page 13-15](#)
- [Default Settings, page 13-20](#)
- [Zone Implementation, page 13-21](#)

Table 9-1 on page 9-4 lists the differences between zones and VSANs.



### Note

For a comprehensive summary of zone implementation, refer to the [“Zone Implementation” section on page 13-21](#).

# Zoning Features

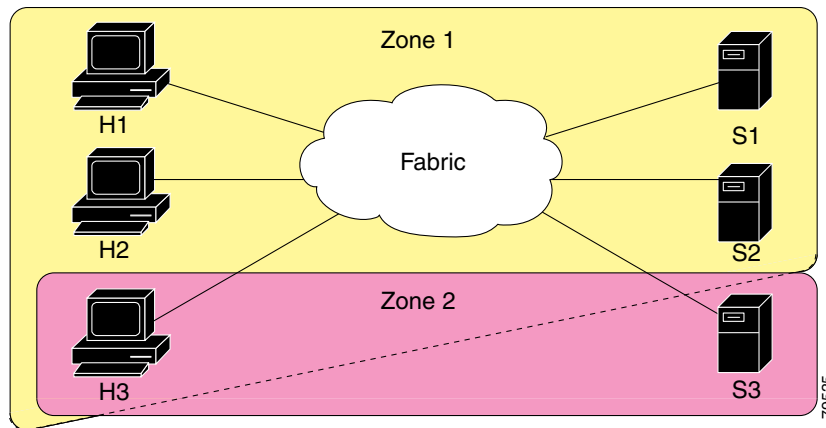
Zoning has the following features:

- A zone consists of multiple zone members.
  - Members in a zone can access each other; members in different zones cannot access each other.
  - If zoning is not activated, all devices are members of the default zone.
  - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
  - Zones can vary in size.
  - Devices can belong to more than one zone.
- A zone set consists of one or more zones.
  - A zone set can be activated or deactivated as a single entity across all switches in the fabric.
  - Only one zone set can be activated at any time.
  - A zone can be a member of more than one zone set.
- Zoning can be administered from any switch in the fabric.
  - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if the option (**zoneset distribute full vsan** command) is enabled in the source switch.
  - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured nondisruptively.
  - New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.
- Zone membership criteria is based on WWNs or FC IDs.
  - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
  - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
  - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
  - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
  - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
  - Domain ID and port number—Specifies the domain ID of an MDS domain and additionally specifies a port belonging to a non-Cisco switch.
  - IP address—Specifies the IP address (and optionally the subnet mask) of an attached device.
- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.

# Zoning Example

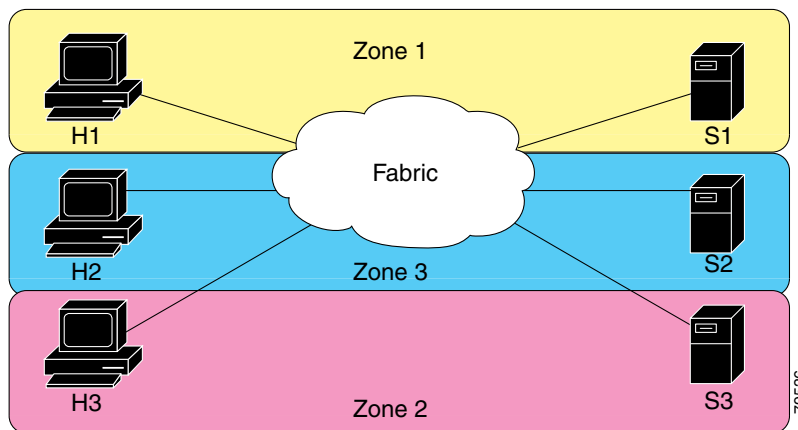
**Figure 13-1** illustrates a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. Note that H3 resides in both zones.

**Figure 13-1 Fabric with Two Zones**



Of course, there are other ways to partition this fabric into zones. **Figure 13-2** illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in zone 1.

**Figure 13-2 Fabric with Three Zones**




# Configuring a Zone

A zone can be configured using one of the following types to assign members:

- pWWN—The WWN of the N or NL port in hex format (for example, 10:00:00:23:45:67:89:ab).
- Fabric port WWN—The WWN of the fabric port name in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID in 0xhhhhhh format (for example, 0xce00d1).
- FC alias—The alias name is in alphabetic characters (for example, Payroll) and denotes a port ID or WWN. The alias can also include multiple members.
- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- IP address—The IP address of an attached device in 32 bytes in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

To configure a zone and assign a zone name, follow these steps:

|                                                                                                                                                           | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Purpose                                                                                                                                                             |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1                                                                                                                                                    | switch# <b>config t</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Enters configuration mode.                                                                                                                                          |
| Step 2                                                                                                                                                    | switch(config)# <b>zone name Zone1 vsan 3</b><br>switch(config-zone)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Configures a zone called Zone 1 for the VSAN called vsan3.                                                                                                          |
| Step 3                                                                                                                                                    | switch(config-zone)# <b>member &lt;type&gt; &lt;value&gt;</b><br>pWWN example:<br>sswitch(config-zone)# <b>member pwwn 10:00:00:23:45:67:89:ab</b><br>Fabric pWWN example:<br>switch(config-zone)# <b>member fwwn 10:01:10:01:10:ab:cd:ef</b><br>FC ID example:<br>switch(config-zone)# <b>member fcid 0xce00d1</b><br>FC alias example:<br>switch(config-zone)# <b>member fcalias Payroll</b><br>Domain ID example:<br>switch(config-zone)# <b>member domain-id 2 portnumber 23</b><br>FC alias example:<br>switch(config-zone)# <b>member ipaddress 10.15.0.0 255.255.0.0</b><br>Local sWWN interface example:<br>switch(config-zone)# <b>member interface fc 2/1</b><br>Remote sWWN interface example:<br>switch(config-zone)# <b>member interface fc2/1 swnn 20:00:00:05:30:00:4a:de</b><br>Domain ID interface example:<br>switch(config-zone)# <b>member interface fc2/1 domain-id 25</b> | Configures a member for the specified zone (Zone1) based on the type (pWWN, fabric pWWN, FC ID, FC alias, domain ID, IP address, or interface) and value specified. |
| <b>Tip</b> Use a relevant display command (for example, <b>show interface</b> or <b>show flogi database</b> ) to obtain the required value in hex format. |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                     |



Note

Interface-based zoning only works with Cisco MDS 9000 family switches. Interface-based zoning does not work if **interop** mode is configured in that VSAN.

**Tip**

Use the **show wwn switch** command to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.

## Configuring Aliases

You can assign an alias name and configure an alias member using either the FC ID, fabric port WWN (fWWN), or pWWN values.

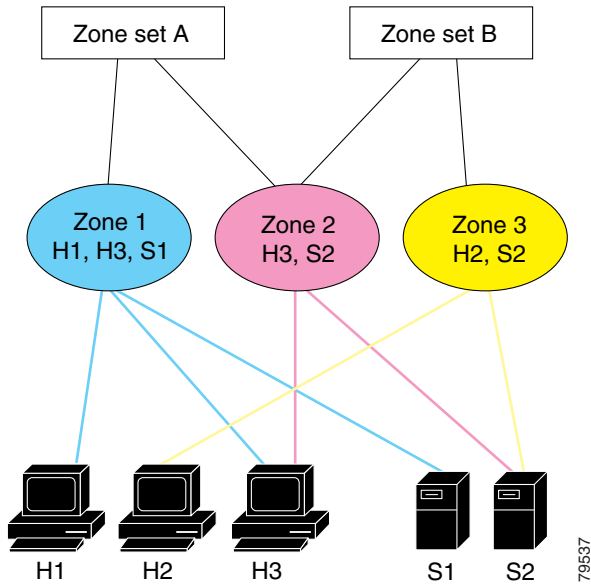
To create an alias using the **fcalias** command, follow these steps:

|               | Command                                                                          | Purpose                                                                                                 |
|---------------|----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b>                                                          | Enters configuration mode.                                                                              |
| <b>Step 2</b> | switch(config)# <b>fcalias name AliasSample vsan 3</b><br>switch-config-fcalias# | Configures an alias name (AliasSample).                                                                 |
| <b>Step 3</b> | switch-config-fcalias# <b>member fcid 0x222222</b>                               | Configures alias members based on the specified FC ID type and value (0x222222).                        |
|               | switch-config-fcalias# <b>member pwnn</b><br>10:00:00:23:45:67:89:ab             | Configures alias members based on the specified port WWN type and value (pWWN 10:00:00:23:45:67:89:ab). |
|               | switch-config-fcalias# <b>member fwnn</b><br>10:01:10:01:10:ab:cd:ef             | Configures alias members based on the specified fWWN type and value (fWWN 10:01:10:01:10:ab:cd:ef).     |
| <b>Note</b>   | Multiple members can be specified on multiple lines.                             |                                                                                                         |

# Zone Sets

In [Figure 13-3](#), two separate sets are created, each with its own membership hierarchy and zone members.

Figure 13-3 Hierarchy of Zone Sets, Zones, and Zone Members



Zones provide a mechanism for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

To create a zone set to include several zones, follow these steps:

|        | Command                                                                                  | Purpose                                                                                                                                                                                                         |
|--------|------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <code>config t</code>                                                            | Enters configuration mode.                                                                                                                                                                                      |
| Step 2 | switch(config)# <code>zoneset name Zoneset1 vsan 3</code><br>switch-config-zoneset#      | Configures a zone set called Zoneset1.<br><b>Note</b> To activate a zone set, you must first create the zone and a zone set.                                                                                    |
| Step 3 | switch-config-zoneset# <code>member Zone1</code><br>switch-config-zoneset#               | Adds Zone1 as a member of the specified zone set (Zoneset1).<br><b>Note</b> If the specified zone name was not previously configured, this command will return the <code>Zone not present</code> error message. |
| Step 4 | switch-config-zoneset# <code>zone name InlineZone1</code><br>switch-config-zoneset-zone# | Adds a zone (InlineZone1) to the specified zone set (Zoneset1).<br><b>Tip</b> Execute this step only if you need to create a zone from a zone set prompt.                                                       |



|                                                                  | Command                                                                                | Purpose                                                                                         |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| Step 5                                                           | switch-config-zoneset-zone# <b>member fcid 0x111112</b><br>switch-config-zoneset-zone# | Adds a new member (FC ID 0x111112) to the newly created zone (InlineZone1).                     |
|                                                                  |                                                                                        | <b>Tip</b> Execute this step only if you need to add a member to a zone from a zone set prompt. |
| <b>Note</b> Multiple members can be specified on multiple lines. |                                                                                        |                                                                                                 |

Zone sets are configured with the names of the member zones. If the zone set is in a configured VSAN, the VSAN is also specified.

## Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone. You can activate a zone set using the **zoneset activate name** command.
- The administrator can modify the full zone set even if a zone set with the same name is active. The changes do not take effect until the zone set is activated with the **zoneset activate name** command.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets. You do not have to issue the **copy running-config startup-config** command to store the active zone set. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. It is not available across switch resets.
- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.

Figure 13-4 shows a zone being added to an activated zone set.

To activate a zone set, follow these steps:

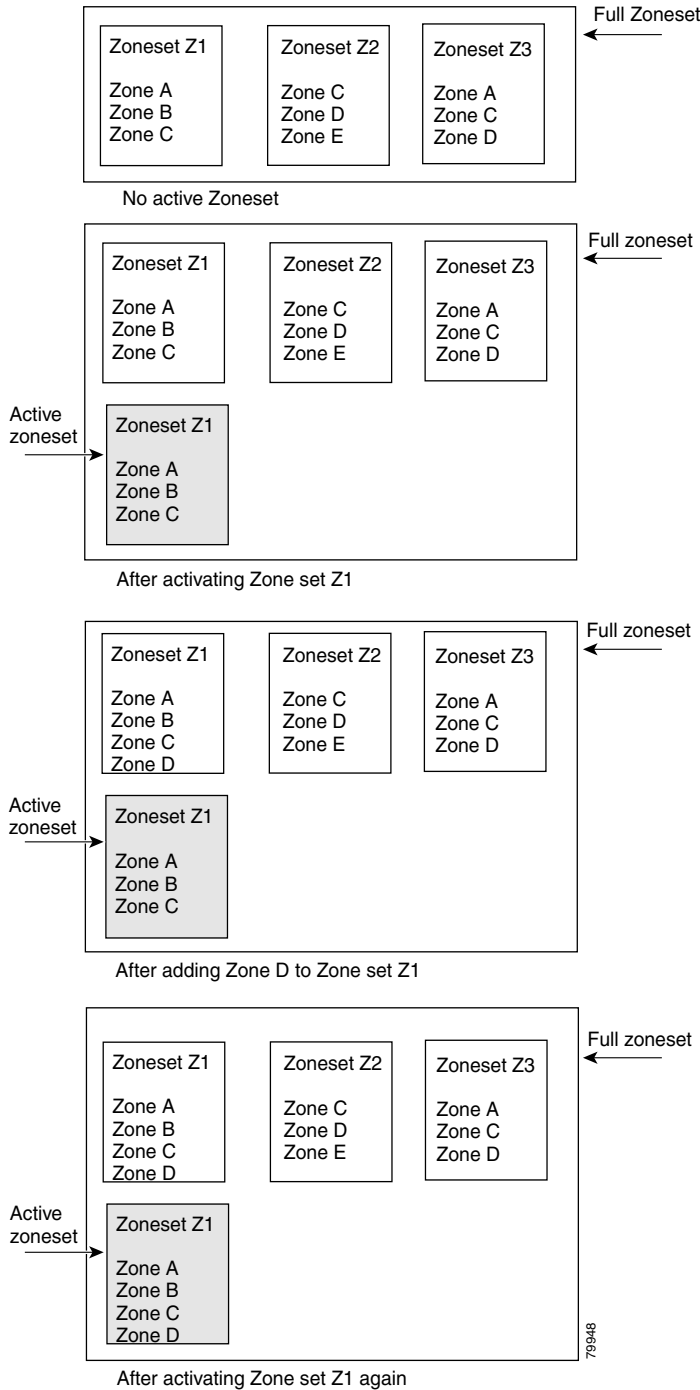
|        | Command                                                         | Purpose                            |
|--------|-----------------------------------------------------------------|------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                      | Enters configuration mode.         |
| Step 2 | switch(config)# <b>zoneset activate name Zoneset1 vsan 3</b>    | Activates the specified zone set.  |
|        | switch(config)# <b>no zoneset activate name Zoneset1 vsan 3</b> | Deactivates the specified zone set |



**Note**

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You don't need to explicitly deactivate the currently active zone set before activating a new zone set.

**Figure 13-4 Active and Full Zone Sets**



## Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

Hard zoning is enforced by the hardware on each frame sent by an Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wire speed. Hard zoning is applied to all forms of zoning.



**Note**

Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

## The Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.



**Note**

Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied amongst members of the default zone. This information is not distributed to all switches; it must be configured in each switch.



**Note**

When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to talk to each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric. The default zone members are explicitly listed when the default policy is configured as **permit** or when a zoneset is active. When the default policy is configured as **deny**, the members of this zone are not explicitly enumerated when you issue the **show zoneset active** command.

To permit or deny traffic in the default zone, follow these steps:

|        | Command                 | Purpose                    |
|--------|-------------------------|----------------------------|
| Step 1 | switch# <b>config t</b> | Enters configuration mode. |

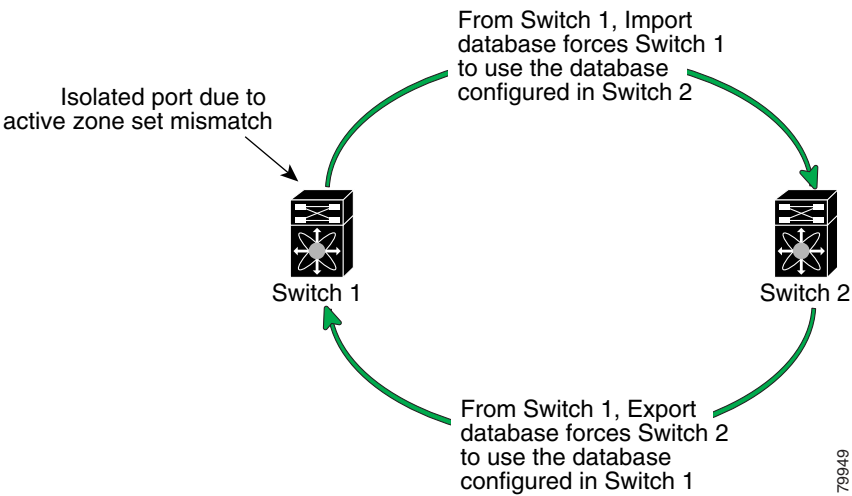
|                                                                                  | Command                                                         | Purpose                                                                     |
|----------------------------------------------------------------------------------|-----------------------------------------------------------------|-----------------------------------------------------------------------------|
| Step 2                                                                           | <code>switch(config)# zone default-zone permit vsan 1</code>    | Permits traffic flow to default zone members.                               |
|                                                                                  | <code>switch(config)# no zone default-zone permit vsan 1</code> | Denies traffic flow to default zone members and reverts to factory default. |
| <b>Note</b> The default settings for default zone configurations can be changed. |                                                                 |                                                                             |

# Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zoneset databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zoneset database and replace the current active zoneset (see [Figure 13-5](#)).
- Export the current database to the neighboring switch (see [Figure 13-5](#)).
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 13-5    Importing and Exporting the Database



Tip

The **import** and **export** commands should be issued from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

To import the zoneset from an adjacent switch, follow these steps:

|        | Command                                              | Purpose                                                                              |
|--------|------------------------------------------------------|--------------------------------------------------------------------------------------|
| Step 1 | switch# <b>zoneset import interface fc1/3 vsan 2</b> | Imports the zoneset from the adjacent switch connected through the VSAN 2 interface. |
|        | switch# <b>zoneset export vsan 5</b>                 | Exports the zoneset to the adjacent switch connected through VSAN 5.                 |

**Note**

You can also issue the **zoneset import** and the **zoneset export** commands for a range of VSANs.

## Distributing Zone Sets

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The **zoneset distribute full vsan** command distributes the full zone set along with the active zone set. The distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

To distribute zone sets, follow these steps:

|        | Command                                                | Purpose                                                        |
|--------|--------------------------------------------------------|----------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                | Enters configuration mode.                                     |
| Step 2 | switch(config)# <b>zoneset distribute full vsan 33</b> | Enables sending a full zone set along with an active zone set. |

## Copying Zone Sets

The active zone set is not a part of the full zone set. You can copy an active zone set using the **zone copy active-zoneset** command. You can not make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated. This command does not distribute zone sets. Since you can not edit an active zone set, this command is helpful in copying an active zone set. You can make a copy and then edit it without altering the existing active zone set. You can copy an active-zone set to a location in bootflash, volatile, slot0, to a remote location (using FTP, SCP, SFTP, or TFTP), or to the full zone set.

**Caution**

Copying an active zone set to a full zone set may overwrite a zone with the same name, if it already exists in the full zone set database.

To copy zone sets, follow this step:

|        | Command                                                                                                  | Purpose                                                          |
|--------|----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| Step 1 | switch# <b>zone copy active-zoneset full-zoneset</b><br>Please enter yes to proceed. (y/n) [n]? <b>y</b> | Makes a copy of the active zone set in the full zone set.        |
|        | switch# <b>zone copy vsan 3 active-zoneset</b><br><b>scp://guest@myserver/tmp/active_zoneset.txt</b>     | Copies the active zone in VSAN 3 to a remote location using SCP. |

# Clearing the Zone Database



## Note

Clearing a zone set only erases the full zone database, not the active zone database.

To clear the zone server database, use the **clear zone database** command.

```
switch# clear zone database vsan 2
```

This command clears all configured information in the zone server for the specified VSAN.



## Note

After issuing a **clear zone database** command, you need to explicitly issue the **copy running-config startup-config** to ensure that the running configuration is used when you next start the switch.

# LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.



## Note

LUN zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2(x) or above.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

Figure 13-6 shows a LUN-based zone example.

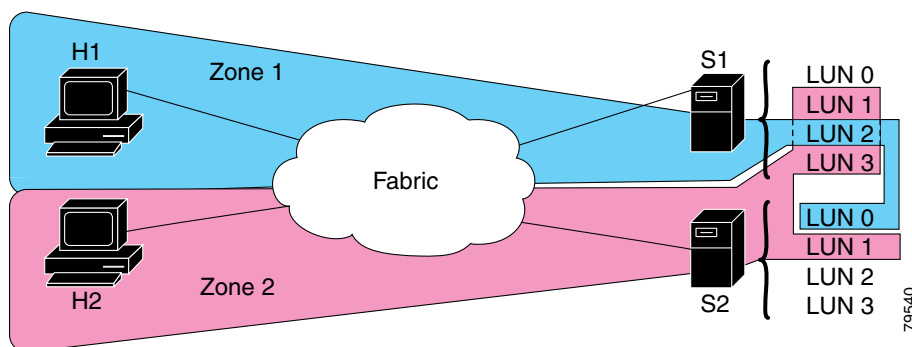
- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUN in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUN in S1 or S2.



## Note

Unzoned LUNs automatically become members of the default zone.

**Figure 13-6 LUN Zoning Access**



To configure a LUN-based zone, follow these steps:

|        | Command                                                                          | Purpose                                                                                                                                            |
|--------|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                       | Enters configuration mode.                                                                                                                         |
| Step 2 | switch(config)# <b>zone name LunSample vsan 2</b><br>switch(config-zone)#        | Configures a zone called LunSample for the specified VSAN (vsan 2).                                                                                |
| Step 3 | switch(config-zone)# <b>member pwwn</b><br><b>10:00:00:23:45:67:89:ab lun 64</b> | Configures a zone member based on the specified pWWN and LUN value.<br><br><b>Note</b> LUN x64 in hex format corresponds to 100 in decimal format. |
|        | switch(config-zone)# <b>member fcid 0x12465</b><br><b>lun 64</b>                 | Configures a zone member based on the FC ID and LUN value.                                                                                         |

**Note**

When LUN 0 is not included within a zone, then, as per standards requirements, control traffic to LUN 0 (for example, REPORT\_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.

**Caution**

LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure **interop** mode in that switch.

## Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each Host Bus Adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided in the preceding section.

**Note**

Refer to the relevant user manuals to obtain the LUN number for each HBA.

**Caution**

If you make any errors when configuring this scenario, you are prone to loose data.

## Read-Only Zoning

**Note**

Read-only zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2(x) or above.

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

You can also configure LUN zones as read-only zones.

## Guidelines to Configure Read-Only Zones

Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

Follow these guidelines when configuring read-only zones:

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, read-only zone has priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure **interop** mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.

The read-only zone feature behaves as designed if FAT16 or FAT32 file system is used with the above-mentioned Windows operating systems.

## Configuring Read-Only Zones

To configure read-only zones, follow these steps:

|        | Command                                                                 | Purpose                                                                                                 |
|--------|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                              | Enters configuration mode.                                                                              |
| Step 2 | switch(config)# <b>zone name Sample2 vsan 2</b><br>switch(config-zone)# | Configures a zone called Sample2 for the specified VSAN (vsan 2).                                       |
| Step 3 | switch123(config-zone)# <b>attribute read-only</b>                      | Sets read-only attributes for the Sample2 zone.<br><b>Note</b> The default is read-write for all zones. |
|        | switch123(config-zone)# <b>no attribute read-only</b>                   | Reverts the Sample2 zone attributes to read-write.                                                      |

To configure the **read-only** option for a default zone, follow these steps:

|        | Command                                                                         | Purpose                                                      |
|--------|---------------------------------------------------------------------------------|--------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                      | Enters configuration mode.                                   |
| Step 2 | switch(config)# <b>zone default-zone vsan 1</b><br>switch(config-default-zone)# | Enters the default-zone submode.                             |
| Step 3 | switch123(config-zone)# <b>attribute read-only</b>                              | Sets read-only attributes for the default zone.              |
|        | switch123(config-zone)# <b>no attribute read-only</b>                           | Reverts the default zone attributes to read-write (default). |



# Viewing Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, alias, or even a keyword like **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed. See Examples 13-1 to 13-14.

## Example 13-1 Displays Zone Information for All VSANs

```
switch# show zone
zone name Zone3 vsan 1
 qwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:9c:48:e5

zone name Zone2 vsan 2
 fwwn 20:41:00:05:30:00:2a:1e
 fwwn 20:42:00:05:30:00:2a:1e
 fwwn 20:43:00:05:30:00:2a:1e

zone name Zone1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1

zone name Techdocs vsan 3
 ip-address 10.15.0.0 255.255.255.0

zone name Zone21 vsan 5
 pwwn 21:00:00:20:37:a6:be:35
 pwwn 21:00:00:20:37:a6:be:39
 fcid 0xe000ef
 fcid 0xe000e0
 symbolic-nodename ign.test
 fwwn 20:1f:00:05:30:00:e5:c6
 fwwn 12:12:11:12:11:12:12:10
 interface fc1/5 swwn 20:00:00:05:30:00:2a:1e
 ip-address 12.2.4.5 255.255.255.0
 fcalias name Alias1 vsan 1
 pwwn 21:00:00:20:37:a6:be:35

zone name Zone2 vsan 11
 interface fc1/5 pwwn 20:4f:00:05:30:00:2a:1e

zone name Zone22 vsan 6
 fcalias name Alias1 vsan 1
 pwwn 21:00:00:20:37:a6:be:35

zone name Zone23 vsan 61
 pwwn 21:00:00:04:cf:fb:3e:7b lun 0000
```

## Example 13-2 Displays Zone Information for a Specific VSAN.

```
switch# show zone vsan 1
zone name Zone3 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:9c:48:e5

zone name Zone2 vsan 1
 fwwn 20:4f:00:05:30:00:2a:1e
 fwwn 20:50:00:05:30:00:2a:1e
```

```

fwwn 20:51:00:05:30:00:2a:1e
fwwn 20:52:00:05:30:00:2a:1e
fwwn 20:53:00:05:30:00:2a:1e

zone name Zone1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1

```

Use the **show zoneset** command to view the configured zone sets.

**Example 13-3 Displays Configured Zone Set Information:**

```

switch# show zoneset vsan 1
zoneset name ZoneSet2 vsan 1
 zone name Zone2 vsan 1
 fwwn 20:4e:00:05:30:00:2a:1e
 fwwn 20:4f:00:05:30:00:2a:1e
 fwwn 20:50:00:05:30:00:2a:1e
 fwwn 20:51:00:05:30:00:2a:1e
 fwwn 20:52:00:05:30:00:2a:1e

 zone name Zone1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1

zoneset name ZoneSet1 vsan 1
 zone name Zone1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1

```

**Example 13-4 Displays Configured Zone Set Information for a Range of VSANs:**

```

switch# show zoneset vsan 2-3
zoneset name ZoneSet2 vsan 2
 zone name Zone2 vsan 2
 fwwn 20:52:00:05:30:00:2a:1e
 fwwn 20:53:00:05:30:00:2a:1e
 fwwn 20:54:00:05:30:00:2a:1e
 fwwn 20:55:00:05:30:00:2a:1e
 fwwn 20:56:00:05:30:00:2a:1e

 zone name Zone1 vsan 2
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1

zoneset name ZoneSet3 vsan 3
 zone name Zone1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1

```

Use the **show zone name** command to display members of a specific zone.

**Example 13-5 Displays Members of a Zone**

```
switch# show zone name Zone1
zone name Zone1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
 pwwn 21:00:00:20:37:9c:48:e5
 fcalias Alias1
```

Use the **show fcalias** command to display fcalias configuration.

**Example 13-6 Displays fcalias Configuration**

```
switch# show fcalias vsan 1
fcalias name Alias2 vsan 1

fcalias name Alias1 vsan 1
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:9c:48:e5
```

Use the **show zone member** command to display all zones to which a member belongs using the FC ID.

**Example 13-7 Displays Membership Status**

```
switch# show zone member pwwn 21:00:00:20:37:9c:48:e5
 VSAN: 1
zone Zone3
zone Zone1
fcalias Alias1
```

Use the **show zone statistics** command to display the number of control frames exchanged with other switches.

**Example 13-8 Displays Zone Statistics**

```
switch# show zone statistics
Statistics For VSAN: 1

Number of Merge Requests Sent: 24
Number of Merge Requests Recvd: 25
Number of Merge Accepts Sent: 25
Number of Merge Accepts Recvd: 25
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
Statistics For VSAN: 2

Number of Merge Requests Sent: 4
Number of Merge Requests Recvd: 4
Number of Merge Accepts Sent: 4
Number of Merge Accepts Recvd: 4
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
```

```

Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0

```

### Example 13-9 Displays LUN Zone Statistics

```

switch# show zone statistics lun-zoning
LUN zoning statistics for VSAN: 1

S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:00:00

Number of Inquiry commands received: 10
Number of Inquiry data No LU sent: 5
Number of Report LUNs commands received: 10
Number of Request Sense commands received: 1
Number of Other commands received: 0
Number of Illegal Request Check Condition sent: 0

S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:00:01

Number of Inquiry commands received: 1
Number of Inquiry data No LU sent: 1
Number of Request Sense commands received: 1
Number of Other commands received: 0
Number of Illegal Request Check Condition sent: 0

```

### Example 13-10 Displays LUN Zone Statistics

```

switch# show zone statistics read-only-zoning
Read-only zoning statistics for VSAN: 2

S-ID: 0x33333, D-ID: 0x11111, LUN: 00:00:00:00:00:00:00:64

Number of Data Protect Check Condition Sent: 12

```

### Example 13-11 Displays Active Zonesets

```

switch# show zoneset active
zoneset name ZoneSet1 vsan 1
 zone name zone1 vsan 1
 fcid 0x080808
 fcid 0x090909
 fcid 0x0a0a0a
 zone name zone2 vsan 1
 * fcid 0xef0000 [pwwn 21:00:00:20:37:6f:db:dd]
 * fcid 0xef0100 [pwwn 21:00:00:20:37:a6:be:2f]

```

### Example 13-12 Displays Brief Descriptions of Zone Sets

```

switch# show zoneset brief
zoneset name ZoneSet1 vsan 1
 zone zone1
 zone zone2

```

### Example 13-13 Displays Active Zones

```

switch# show zone active
zone name zone1 vsan 1
 fcid 0x080808
 fcid 0x090909

```

```

fcid 0x0a0a0a

zone name zone2 vsan 1
* fcid 0xef0000 [pwwn 21:00:00:20:37:6f:db:dd]
* fcid 0xef0100 [pwwn 21:00:00:20:37:a6:be:2f]

```

### Example 13-14 Displays Zone Status

```

switch# show zone status
VSAN: 1 default-zone: deny distribute: full Interop: Off
Full Zoning Database :
 Zonesets:1 Zones:11 Aliases:0
Active Zoning Database :
 Name: zoneset-1 Zonesets:1 Zones:11 Aliases:0
Status: Activation completed at Thu Feb 13 10:22:34 2003

VSAN: 2 default-zone: deny distribute: full Interop: Off
Full Zoning Database :
 Zonesets:1 Zones:10 Aliases:0
Active Zoning Database :
 Name: zoneset-2 Zonesets:1 Zones:10 Aliases:0
Status: Activation completed at Thu Feb 13 10:23:12 2003

VSAN: 3 default-zone: deny distribute: full Interop: Off
Full Zoning Database :
 Zonesets:1 Zones:10 Aliases:0
Active Zoning Database :
 Name: zoneset-3 Zonesets:1 Zones:10 Aliases:0
Status: Activation completed at Thu Feb 13 10:23:50 2003

```

Use the **show zone** command to display the zone attributes for all configured zones.

### Example 13-15 Displays Zone Statistics

```

switch123(config-zone)# do show zone
zone name lunSample vsan 1 <-----Read-write attribute
zone name ReadOnlyZone vsan 2
 attribute read-only <-----Read-only attribute

```

Use the **show running** and **show zone active** commands display the configured interface-based zones (see [Example 13-16](#) and [Example 13-17](#)).

### Example 13-16 Displays the Interface-Based Zones

```

switch# show running
zone name if-zone vsan 1
 member interface fc2/15 swrn 20:00:00:0c:88:00:4a:e2
 member fwwn 20:4f:00:0c:88:00:4a:e2
 member interface fc2/1 swrn 20:00:00:05:30:00:4a:9e
 member pwwn 22:00:00:20:37:39:6b:dd

```

### Example 13-17 Displays the fWWNs and Interfaces in a Active Zone

```

switch# show zone active
zone name if-zone vsan 1
* fcid 0x7e00b3 [interface fc2/15 swrn 20:00:00:0c:88:00:4a:e2]
* fcid 0x7e00b1 [interface fc2/15 swrn 20:00:00:0c:88:00:4a:e2]
* fcid 0x7e00ac [interface fc2/15 swrn 20:00:00:0c:88:00:4a:e2]
* fcid 0x7e00b3 [fwwn 20:4f:00:0c:88:00:4a:e2]
* fcid 0x7e00b1 [fwwn 20:4f:00:0c:88:00:4a:e2]
* fcid 0x7e00ac [fwwn 20:4f:00:0c:88:00:4a:e2]

```

```
interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
```

A similar output is also available on the remote switch (see [Example 13-18](#)).

**Example 13-18 Displays the Local Interface Active Zone Details for a Remote Switch**

```
switch# show zone active
zone name if-zone vsan 1
 * fcid 0x7e00b3 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
 * fcid 0x7e00b1 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
 * fcid 0x7e00ac [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
 * fcid 0x7e00b3 [fwwn 20:4f:00:0c:88:00:4a:e2]
 * fcid 0x7e00b1 [fwwn 20:4f:00:0c:88:00:4a:e2]
 * fcid 0x7e00ac [fwwn 20:4f:00:0c:88:00:4a:e2]
 interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
```

# Default Settings

[Table 13-1](#) lists the default settings for zone parameters.

**Table 13-1    Default Zone Parameters**

| Parameters               | Default                                  |
|--------------------------|------------------------------------------|
| Default zone policy      | Denied to all members.                   |
| Full zone set distribute | The full zone set(s) is not distributed. |
| Read-only zones          | Read-write attributes for all zones.     |

# Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be changed (more secure).
- Name server queries are soft-zoned.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zonesets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic

You can additionally configure the following zone features if required:

- Propagate full zone sets to all switches on a per VSAN basis using the **zoneset distribute full-database vsan** command.
- Change the default policy for unzoned members using the **zone default permit vsan** command.
- Inter-operate with other vendors by configuring a VSAN in the interop mode using the **vsan 1 interop** command.
  - Configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other
- Bring E ports out of isolation using the **export** or **import** commands







## Configuring Inter-VSAN Routing

---

This chapter explains the Inter-VSAN Routing (IVR) feature and provides details on sharing resources across VSANs using IVR management interfaces provided in the switch.

This chapter includes the following sections:

- [About IVR, page 14-2](#)
- [IVR Features, page 14-2](#)
- [IVR Terminology, page 14-3](#)
- [IVR Guidelines, page 14-4](#)
- [Configuring IVR, page 14-5](#)
- [Unique Domain ID Configuration Options, page 14-5](#)
- [Enabling IVR, page 14-6](#)
- [Configuring an IVR Topology, page 14-6](#)
- [Creating IVZs and IVZSs, page 14-8](#)
- [IVR Interoperability, page 14-12](#)
- [IVR Using LUN Zoning or Read-Only Zoning, page 14-12](#)
- [Clearing the IVZ Database, page 14-12](#)
- [Specifying IVR logging Levels, page 14-13](#)
- [Viewing IVR Information, page 14-13](#)

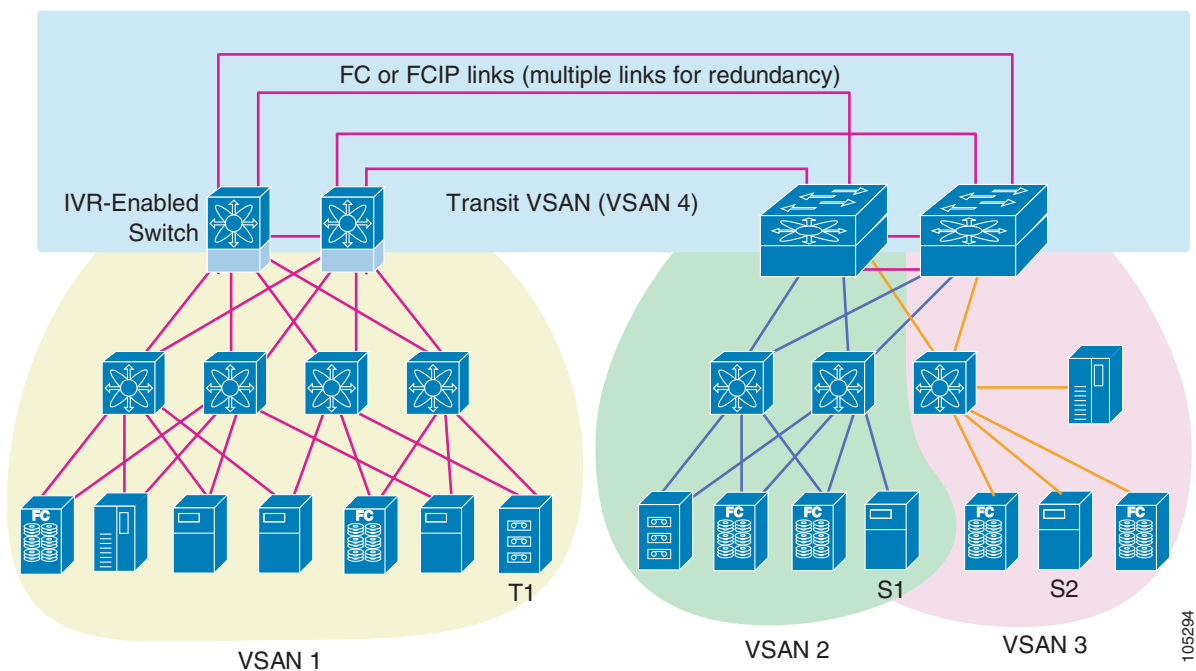
## About IVR

Virtual SANs (VSANs) improve Storage Area Network (SAN) scalability, availability, and security by allowing multiple Fibre Channel SANs to share a common physical infrastructure of switches and ISLs. These benefits are derived from the separation of Fibre Channel services in each VSAN and isolation of traffic between VSANs. Data traffic isolation between the VSANs also inherently prevents sharing of resources attached to a VSAN, like robotic tape libraries. Using IVR, resources across VSANs are accessed without compromising other VSAN benefits.

Data traffic is transported between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric. FC control traffic does not flow between VSANs, nor can initiators access any resource across VSANs aside from the designated ones. Valuable resources like tape libraries are easily shared across VSANs without compromise.

IVR is not limited to VSANs present on a common switch. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions (see [Figure 14-1](#)). See the “[Sample Configuration](#)” section on page 14-16 to configure this example.

**Figure 14-1 Traffic Continuity Using IVR and FCIP**



## IVR Features

IVR has the following features:

- Accesses resources across VSANs without compromising other VSAN benefits.
- Transports data traffic between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric.
- Shares valuable resources (like tape libraries) across VSANs without compromise.

- Provides efficient business continuity or disaster recovery solutions when used in conjunction with FCIP.
- Compliant with Fibre Channel standards compliant.
- Incorporates with third-party switches—if the IVR-enabled VSANs are configured in **interop 1** mode.

## IVR Terminology

The terms used in this chapter are explained in this section.

- *Native VSAN*  
The VSAN to which an end device logs on is called a native VSAN for that end device.
- *Inter-VSAN Zone (IVZ)*  
Defines a set of end devices that are allowed to communicate across VSANs within their interconnected SAN fabric. This definition is based on their port World Wide Names (pWWNs) and their native VSAN association. You can configure up to 200 IVZs and 2000 IVZ members on any switch in the Cisco MDS 9000 Family.
- *Inter-VSAN ZoneSets (IVZS)*  
One or more IVZs make up an IVZS. You can configure up to 32 IVZSs on any switch in the Cisco MDS 9000 Family. Only one IVZS can be active at any time.
- *IVR Path*  
An IVR path is a set of switches and inter-switch links via which a frame from one end-device in one VSAN can reach another end-device in some other VSAN. Multiple paths can exist between two such end-devices.
- *IVR-Enabled Switch*  
A switch in which the IVR feature is enabled.
- *Edge VSAN*  
An edge VSAN refers to a VSAN which initiates (source edge-VSAN) or terminates (destination edge-VSAN) an IVR path. Edge VSANs may be adjacent to each other or they may be connected by one or more transit VSANs. In [Figure 14-1](#), VSANs 1, 2, and 3 are edge VSANs.




---

**Note** An edge VSAN for one IVR path can be a transit VSAN for another IVR path.

---

- *Transit VSAN*  
Transit VSAN is a VSAN that exists along an IVR path from the source edge VSAN of that path to the destination edge VSAN of that path. In [Figure 14-1](#), VSAN 4 is a transit VSAN.




---

**Note** When the source and destination edge VSANs are adjacent to each other, then a transit VSAN is not required between them.

---

- *Border Switch*  
An IVR-enabled switch that is a member of two or more VSANs, as identified in [Figure 14-1](#).
- *Edge Switch*  
A switch to which a member of an IVR zone has logged in. Edge switches are oblivious to the IVR configurations in the border switches. Edge switches need not be IVR enabled.

# IVR Guidelines

Before configuring an IVR SAN fabric, consider the following guidelines:

- Configure unique domain IDs across all VSANs and switches participating in IVR operations. The following switches participate in IVR operations:
  - All edge switches in the edge VSANs (source and destination)
  - All switches in transit VSANs
- Configure IVR only in the relevant border switches.
- An Enterprise License Package is required for this feature.



**Tip**

If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.



**Note**

IVR-enabled VSANs must be configured in **no interop** (default) mode or **interop 1** mode.

## Domain ID Guidelines

Domain IDs must be unique across inter-connected VSANs. To ensure unique domain IDs across inter-connected VSAN, follow these guidelines:

- Minimize the number of switches that require a domain ID assignment. This ensures minimum traffic disruption.
- Minimize the coordination between interconnected VSANs, when configuring the SAN for the first time, as well as when you add each new switch.

## Transit VSANs Guidelines

Consider the following guidelines for transit VSANs:

- Besides defining the IVZ membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVZ overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
  - If two edge VSANs in an IVZ do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVZ will not overlap if IVR is not enabled on a switch that is a member of both the source and destinations edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVZs. Sometimes, a transit VSAN can also double-up as an edge VSAN in another IVZ.

## Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require SAN-OS Release 1.3(1) or higher.
- A border switch must be a member of two or more VSANs.
- A border switch that facilitates IVR communications must be IVR enabled.
- To provide redundant paths between active IVZ members, IVR can (optionally) be enabled on additional border switches.
- The VSAN topology configuration must be updated before a border switch is added or removed.

## Configuring IVR

To configure IVR in a SAN fabric, follow these steps.

- 
- |               |                                                                                                                                                                 |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Verify that unique domain IDs are configured in all switches and VSANs participating in IVR (see <a href="#">Chapter 24, “Configuring Domain Parameters”</a> ). |
| <b>Step 2</b> | Enable IVR in the border switches (see <a href="#">Enabling IVR, page 14-6</a> ).                                                                               |
| <b>Step 3</b> | Create and activate the required IVR topology in <i>all</i> the IVR-enabled border switches (see <a href="#">Configuring an IVR Topology, page 14-6</a> ).      |
| <b>Step 4</b> | Create and activate IVZSs in <i>all</i> the IVR-enabled border switches (see the <a href="#">Creating IVZs and IVZSs, page 14-8</a> ).                          |
| <b>Step 5</b> | Verify the IVR configuration (see <a href="#">Viewing IVR Information, page 14-13</a> ).                                                                        |
- 

## Unique Domain ID Configuration Options

You can configure domain IDs using one of two options:

- Configure allowed-domains list using the Domain Manager MIBs so that the domains in different VSANs are non-overlapping on all participating switches and VSANs.
- Configure static, non-overlapping domains (using the CLI) for each participating switch and VSAN (see [Chapter 24, “Configuring Domain Parameters”](#)).

## Enabling IVR

The IVR feature must be enabled in all border switches in the fabric that participate in the IVR. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. To begin configuring the IVR feature, you must explicitly enable IVR on the required switches in the fabric.

The configuration and verification commands for the IVR feature are only available when IVR is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable IVR on any participating switch, follow these steps:

|        | Command                              | Purpose                                |
|--------|--------------------------------------|----------------------------------------|
| Step 1 | switch# <b>config t</b>              | Enters configuration mode.             |
| Step 2 | switch(config)# <b>ivr enable</b>    | Enables IVR on that switch.            |
|        | switch(config)# <b>no ivr enable</b> | Disables (default) IVR on that switch. |

## Configuring an IVR Topology

This section explains the process used to create, activate, and clear an IVR topology.

### Creating an IVR Topology

You must create the IVR topology in every IVR-enabled switch in the fabric. You can have up to 64 VSANs in an IVR topology. Specify the IVR topology using the following information:

- The switch WWNs of the IVR-enabled switches (use the **show wwn switch** command).
- A minimum of two VSANs to which the IVR-enabled switch belongs.
- The autonomous fabric ID (AF ID) which distinguishes two VSANs that are logically and physically separate, but have the same VSAN number. SAN-OS Release 1.3(1) supports only one AF ID.



**Note** The use of a single AF ID does not allow for segmented VSANs in an inter-VSAN topology.

To create an IVR topology, follow these steps:

|        | Command                                                                                                                  | Purpose                                                                   |
|--------|--------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                  | Enters configuration mode.                                                |
| Step 2 | switch(config)# <b>ivr vsan-topology database</b><br>switch(config-ivr-topology-db)#                                     | Enters the VSAN topology database configuration mode for the IVR feature. |
| Step 3 | switch(config-ivr-topology-db)# <b>autonomous-fabric-id 1</b><br><b>switch 20:00:00:05:30:01:1b:b8 vsan-ranges 1-2,6</b> | Configures VSANs 1, 2, and 6 to participate in IVR in this switch.        |
| Step 4 | switch(config-ivr-topology-db)# <b>autonomous-fabric-id 1</b><br><b>switch 20:00:00:05:30:01:1b:c2 vsan-ranges 1-2</b>   | Configures VSANs 1 and 2 to participate in IVR in this switch.            |
| Step 5 | switch(config-ivr-topology-db)# <b>end</b><br>switch#                                                                    | Reverts to EXEC mode.                                                     |

View your configured IVR topology using the **show ivr vsan-topology** command.

```
switch# show ivr vsan-topology
```

| AFID | SWITCH WWN                | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1    | 20:00:00:05:30:01:1b:c2 * | no     | yes 1-2    |
| 1    | 20:02:00:44:22:00:4a:05   | no     | yes 1-2,6  |
| 1    | 20:02:00:44:22:00:4a:07   | no     | yes 2-5    |

Total: 3 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is INACTIVE



#### Note

Ensure to repeat this configuration in all IVR-enabled switches.



#### Tip

Transit VSANs are deduced based on your configuration. The IVR feature does not have an explicit transit-VSAN configuration. In the example used above, VSAN 2 is the transit VSAN between VSANs 1 and 3.

## Activating an IVR Topology

After configuring the IVR topology, you must activate it.

To activate a configured IVR topology, follow these steps:

|        | Command                                           | Purpose                                |
|--------|---------------------------------------------------|----------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#        | Enters configuration mode.             |
| Step 2 | switch(config)# <b>ivr vsan-topology activate</b> | Activates the configured IVR topology. |



#### Caution

Active IVR topologies cannot be deactivated.

View your active IVR topology using the **show ivr vsan-topology** command.

```
switch# show ivr vsan-topology
```

| AFID | SWITCH WWN                | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1    | 20:00:00:05:30:01:1b:c2 * | yes    | yes 1-2    |
| 1    | 20:02:00:44:22:00:4a:05   | yes    | yes 1-2,6  |
| 1    | 20:02:00:44:22:00:4a:07   | yes    | yes 2-5    |

Total: 3 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is **ACTIVE**

Last activation time: Mon Mar 24 07:19:53 1980

## Clearing the IVR Topology

You can clear a configured IVR topology using the **no ivr vsan-topology database** command in configuration mode.



### Note

You can replace the active IVR topology with another IVR topology configuration by issuing the **ivr vsan-topology activate** command. Active IVR topologies cannot be deactivated.

To clear a previously-created IVR topology, follow these steps:

|        | Command                                              | Purpose                                     |
|--------|------------------------------------------------------|---------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#           | Enters configuration mode.                  |
| Step 2 | switch(config)# <b>no ivr vsan-topology database</b> | Clears the previously-created IVR topology. |

## Creating IVZs and IVZSs

As part of the IVR configuration, you need to configure one or more IVZs to enable cross-VSAN communication. To achieve this result, you must specify each IVZ as a set of (pWWN, VSAN) entries. Like zones, several IVZs can be configured to belong to an IVRS. You can define several IVZSs and activate only one of the defined IVZSs.



### Note

The same IVZS must be activated on *all* the IVR-enabled switches.

## Zones versus IVZs

Table 14-1 identifies the key differences between IVZs and Zones.

**Table 14-1 Key Differences between IVZs and Zones**

| IVZs                                                             | Zones                                                                        |
|------------------------------------------------------------------|------------------------------------------------------------------------------|
| IVZ membership is specified using the VSAN and pWWN combination. | Zone membership is specified using pWWN, fabric WWN, sWWN, or the fabric ID. |
| Default zone policy is always deny (not configurable).           | Default zone policy is deny (configurable).                                  |

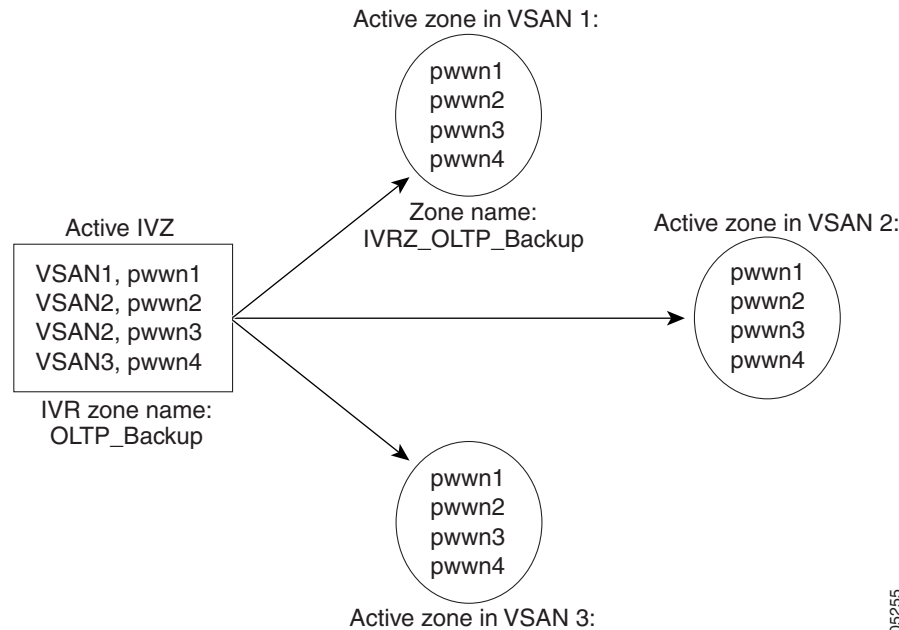


## Automatic IVZ Creation

Figure 14-2 depicts an IVZ consisting of four members. To allow *pwwn1* to communicate with *pwwn2*, they must be in the same zone in VSAN 1, as well as in VSAN 2. If they are not in the same zone, then the hard-zoning ACL entries will prohibit *pwwn1* from communicating with *pwwn2*.

A zone corresponding to each active IVZ is automatically created in each edge VSAN specified in the active IVZ. All pWWNs in the IVZ are members of these zones in each VSAN.

**Figure 14-2 Creating Zones on IVZ Activation**



The zones are created automatically by the IVR process when an IVZS is activated. They are not stored in full zoneset database and are lost when the switch reboots or when a new zoneset is activated. The IVR feature monitors these events and adds the zones corresponding to the active IVZS configuration when a new zoneset is activated. Like zonesets, IVR zonesets are also activated non-disruptively.



### Note

If pwwn1 and pwwn2 are in an IVZ in the current as well as the new IVZS then activation of the new IVZS does not cause any traffic disruption between them.

## Configuring and Activating IVZs and IVZSs

IVZ and IVZS names are restricted to 64 alphanumeric characters.

To configure IVZs and IVZSs, follow these steps:

|        | Command                                                                      | Purpose                                |
|--------|------------------------------------------------------------------------------|----------------------------------------|
| Step 1 | switch# <b>config t</b>                                                      | Enters configuration mode.             |
| Step 2 | switch(config)# <b>ivr zone name Ivz_vsan2-3</b><br>switch(config-ivr-zone)# | Creates an IVR zone named Ivz_vsan2-3. |

|         | Command                                                                                                | Purpose                                             |
|---------|--------------------------------------------------------------------------------------------------------|-----------------------------------------------------|
| Step 3  | <code>switch(config-ivr-zone)# member pwwn 21:00:00:e0:8b:02:ca:4a vsan 3</code>                       | Adds the specified pWWN in VSAN 3 as an IVZ member. |
| Step 4  | <code>switch(config-ivr-zone)# member pwwn 21:00:00:20:37:c8:5c:6b vsan 2</code>                       | Adds the specified pWWN in VSAN 2 as an IVZ member. |
| Step 5  | <code>switch(config-ivr-zone)# exit</code><br><code>switch(config)#</code>                             | Reverts to configuration mode.                      |
| Step 6  | <code>switch(config)# ivr zone name Ivz_vsan4-5</code><br><code>switch(config-ivr-zone)#</code>        | Creates an IVR zone named Ivz_vsan2-3.              |
| Step 7  | <code>switch(config-ivr-zone)# member pwwn 21:00:00:e0:8b:06:d9:1d vsan 4</code>                       | Adds the specified pWWN in VSAN 4 as an IVZ member. |
| Step 8  | <code>switch(config-ivr-zone)# member pwwn 21:01:00:e0:8b:2e:80:93 vsan 4</code>                       | Adds the specified pWWN in VSAN 4 as an IVZ member. |
| Step 9  | <code>switch(config-ivr-zone)# member pwwn 10:00:00:00:c9:2d:5a:dd vsan 5</code>                       | Adds the specified pWWN in VSAN 5 as an IVZ member. |
| Step 10 | <code>switch(config-ivr-zone)# exit</code><br><code>switch(config)#</code>                             | Reverts to configuration mode.                      |
| Step 11 | <code>switch(config)# ivr zoneset name Ivz_zoneset1</code><br><code>switch(config-ivr-zoneset)#</code> | Creates an IVR zoneset named Ivz_zoneset1.          |
| Step 12 | <code>switch(config-ivr-zoneset)# member Ivz_vsan2-3</code>                                            | Adds the Ivz_vsan2-3 IVZ as an IVZS member.         |
| Step 13 | <code>switch(config-ivr-zoneset)# member Ivz_vsan4-5</code>                                            | Adds the Ivz_vsan4-5 IVZ as an IVZS member.         |
| Step 14 | <code>switch(config-ivr-zoneset)# exit</code><br><code>switch(config)#</code>                          | Returns to configuration mode.                      |
| Step 15 | <code>switch(config)# ivr zoneset activate name IVR_ZoneSet1</code>                                    | Activates the newly-created IVZS.                   |
|         | <code>switch(config)# ivr zoneset activate name IVR_ZoneSet1 force</code>                              | Forcefully activates the specified IVZS.            |
|         | <code>switch(config)# no ivr zoneset activate name IVR_ZoneSet1</code>                                 | Deactivates the specified IVZS.                     |
| Step 16 | <code>switch(config-ivr-zoneset)# end</code><br><code>switch#</code>                                   | Returns to EXEC mode.                               |

## Using the force Option

Use the **force** option to activate the specified IVZS. [Table 14-2](#) lists the various scenarios with and without the force option...

**Table 14-2** *IVR Scenarios with and without the force Option.*

| Case | Default Zone Policy | Active ZoneSet before IVRZ Activation               | Force Option Used? | IVZS Activation Status | Active IVRZ Created? | Possible Traffic Disruption |
|------|---------------------|-----------------------------------------------------|--------------------|------------------------|----------------------|-----------------------------|
| 1    | Deny                | No active zone set                                  | No                 | Failure                | No                   | No                          |
| 2    |                     |                                                     | Yes                | Success                | Yes                  | No                          |
| 3    | Deny                | Active zone set present                             | No/Yes             | Success                | Yes                  | No                          |
| 4    | Permit              | No active zone set<br>or<br>Active zone set present | No                 | Failure                | No                   | No                          |
| 5    |                     |                                                     | Yes                | Success                | Yes                  | Yes                         |



### Caution

Using the **force** option of IVZS activation may cause traffic disruption, even for devices that are not involved in IVR. For example, if your configuration does not have any active zone sets and the default zone policy is permit, then an IVZS activation will fail. However, IVZS activation will go through if the **force** option is used. Since zones are created in the edge VSANs corresponding to each IVZ, traffic may be disrupted in edge VSANs where the default zone policy is permit.



### Tip

We recommend you use the Case 3 scenario provided in [Table 14-2](#).

View your IVZ configuration using the **show ivr zone** command.

```
switch# show ivr zone

zone name Ivz_vsan2-3
 pwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwn 21:00:00:20:37:c8:5c:6b vsan 2

zone name Ivz_vsan4-5
 pwn 21:00:00:e0:8b:06:d9:1d vsan 4
 pwn 21:01:00:e0:8b:2e:80:93 vsan 4
 pwn 10:00:00:00:c9:2d:5a:dd vsan 5
```

View your IVZS configuration using the **show ivr zoneset** command.

```
switch# show ivr zoneset

zoneset name ivr_qa_zs_all
 zone name Ivz_vsan2-3
 pwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwn 21:00:00:20:37:c8:5c:6b vsan 2

 zone name Ivz_vsan4-5
 pwn 21:00:00:e0:8b:06:d9:1d vsan 4
 pwn 21:01:00:e0:8b:2e:80:93 vsan 4
 pwn 10:00:00:00:c9:2d:5a:dd vsan 5
```

View your active IVZS configuration status using the **show ivr zoneset active** command.

```
switch# show ivr zoneset active

zoneset name ivr_qa_zs_all
 zone name Ivz_vsan2-3
 * pwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwn 21:00:00:20:37:c8:5c:6b vsan 2

 zone name Ivz_vsan4-5
 pwn 21:00:00:e0:8b:06:d9:1d vsan 4
 * pwn 21:01:00:e0:8b:2e:80:93 vsan 4
 pwn 10:00:00:00:c9:2d:5a:dd vsan 5
```



**Tip**

Ensure to repeat this configuration in all border switches participating in the IVR configuration.



**Note**

Using the Cisco MDS Fabric Manager, you can distribute IVZ configurations to all IVR-capable switches in the interconnected VSAN network. Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

## IVR Interoperability

When using the IVR feature, all border switches in a given fabric must be Cisco MDS switches. However, other switches in the fabric may be non-MDS switches. For example, end devices that are members of the active IVZS may be connected to non-MDS switches. Non-MDS switches may also be present in the transit VSAN(s) or in the edge (VSANs) if the **interop-mode 1** option is enabled (see the [“Configuring the Switch for Interoperability”](#) section on page 29-21).

## IVR Using LUN Zoning or Read-Only Zoning

LUN-zoning and read-only zoning can be used between members of active IVR zones. To configure this service, you need to create and activate LUN-zones and/or read-only zones between the desired IVZ members in all relevant edge VSANs using the zoning interface.

The LUN zoning and read-only zoning features cannot be configured in a IVZS setup.

## Clearing the IVZ Database



**Note**

Clearing a zoneset only erases the configured zone database, not the active zone database.

To clear the IVZ database, use the **clear ivr zone database** command.

```
switch# clear ivr zone database
```

This command clears all configured IVZ information.

**Note**

After issuing a **clear ivr zone database** command, you need to explicitly issue the **copy running-config startup-config** to ensure that the running configuration is used when you next start the switch.

## Specifying IVR logging Levels

To configure the severity level for logging messages from the IVR feature, follow these steps:

|        | Command                                    | Purpose                                                                                                                                                         |
|--------|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)# | Enters configuration mode.                                                                                                                                      |
| Step 2 | switch(config)# <b>logging level ivr 4</b> | Configures Telnet or SSH logging for the IVR feature at level 4 (warning). As a result, logging messages with a severity level of 4 or above will be displayed. |

View the configured logging level for the IVR feature using the **show logging level** command.

```
switch# show logging level
Facility Default Severity Current Session Severity

...
ivr 5 4
...
0(emergencies) 1(alerts) 2(critical)
3(errors) 4(warnings) 5(notifications)
6(information) 7(debugging)
```

## Viewing IVR Information

You can view IVR information by using the **show ivr** set of commands. If you request information for a specific object (for example, a specific zone, zoneset, VSAN, alias, or even a keyword like **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed. See Examples 14-1 to 14-12.

### Example 14-1 Displays the Configured IVR VSAN Topology

```
switch# show ivr vsan-topology
AFID SWITCH WWN Active Cfg. VSANS

1 20:00:00:05:30:01:1b:c2 * yes yes 1-2
1 20:02:00:44:22:00:4a:05 yes yes 1-2,6
1 20:02:00:44:22:00:4a:07 yes yes 2-5
```

Total: 5 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE  
Last activation time: Sat Mar 22 21:46:15 1980

The \* indicates the local switch.

**Example 14-2 Displays the Active IVR VSAN Topology**

```
switch# show ivr vsan-topology active
AFID SWITCH WNN Active Cfg. VSANS

 1 20:00:00:05:30:01:1b:c2 * yes yes 1-2
 1 20:02:00:44:22:00:4a:05 yes yes 1-2,6
 1 20:02:00:44:22:00:4a:07 yes yes 2-5

Total: 5 entries in active IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE
Last activation time: Sat Mar 22 21:46:15
```

**Example 14-3 Displays the Configured IVR VSAN Topology**

```
switch# show ivr vsan-topology configured
AFID SWITCH WNN Active Cfg. VSANS

 1 20:00:00:05:30:01:1b:c2 * yes yes 1-2
 1 20:02:00:44:22:00:4a:05 yes yes 1-2,6
 1 20:02:00:44:22:00:4a:07 yes yes 2-5

Total: 5 entries in configured IVR VSAN-Topology
```

**Example 14-4 Displays the IVZ Configuration**

```
switch# show ivr zone
zone name Ivz_vsan2-3
 pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwwn 21:00:00:20:37:c8:5c:6b vsan 2

zone name ivr_qa_z_all
 pwwn 21:00:00:e0:8b:06:d9:1d vsan 1
 pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
 pwwn 10:00:00:00:c9:2d:5a:dd vsan 1
 pwwn 10:00:00:00:c9:2d:5a:de vsan 2
 pwwn 21:00:00:20:37:5b:ce:af vsan 6
 pwwn 21:00:00:20:37:39:6b:dd vsan 6
 pwwn 22:00:00:20:37:39:6b:dd vsan 3
 pwwn 22:00:00:20:37:5b:ce:af vsan 3
 pwwn 50:06:04:82:bc:01:c3:84 vsan 5
```

**Example 14-5 Displays the Active IVZS Configuration**

```
switch# show ivr zoneset active
zoneset name IVR_ZoneSet1
 zone name Ivz_vsan2-3
 pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

**Example 14-6 Displays Information for a Specified IVZ**

```
switch# show ivr zone name Ivz_vsan2-3
zone name Ivz_vsan2-3
 pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

**Example 14-7 Displays the Specified Zone in the Active IVZS**

```
switch# show ivr zone name Ivz_vsan2-3 active
```

```

zone name Ivz_vsan2-3
 pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwwn 21:00:00:20:37:c8:5c:6b vsan 2

```

#### Example 14-8 Displays the IVZS Configuration

```

switch# show ivr zoneset
zoneset name ivr_qa_zs_all
 zone name ivr_qa_z_all
 pwwn 21:00:00:e0:8b:06:d9:1d vsan 1
 pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
 pwwn 10:00:00:00:c9:2d:5a:dd vsan 1
 pwwn 10:00:00:00:c9:2d:5a:de vsan 2
 pwwn 21:00:00:20:37:5b:ce:af vsan 6
 pwwn 21:00:00:20:37:39:6b:dd vsan 6
 pwwn 22:00:00:20:37:39:6b:dd vsan 3
 pwwn 22:00:00:20:37:5b:ce:af vsan 3
 pwwn 50:06:04:82:bc:01:c3:84 vsan 5

zoneset name IVR_ZoneSet1
 zone name Ivz_vsan2-3
 pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwwn 21:00:00:20:37:c8:5c:6b vsan 2

```

#### Example 14-9 Displays Brief Information for an IVR VSAN Topology

```

switch# show ivr zoneset brief Active
zoneset name IVR_ZoneSet1
 zone name Ivz_vsan2-3

```

#### Example 14-10 Displays Brief Information for the Active IVZS

```

switch# show ivr zoneset brief Active
zoneset name IVR_ZoneSet1
 zone name Ivz_vsan2-3

```

#### Example 14-11 Displays Status Information for the IVZ

```

switch# show ivr zoneset status
Zoneset Status

```

|                    |                            |
|--------------------|----------------------------|
| name               | : IVR_ZoneSet1             |
| state              | : activation success       |
| last activate time | : Sat Mar 22 21:38:46 1980 |
| force option       | : off                      |

```

status per vsan:

```

| vsan | status |
|------|--------|
| 1    | active |
| 2    | active |

#### Example 14-12 Displays the Specified Zoneset

```

switch# show ivr zoneset name IVR_ZoneSet1
zoneset name IVR_ZoneSet1
 zone name Ivz_vsan2-3
 pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
 pwwn 21:00:00:20:37:c8:5c:6b vsan 2

```

# Sample Configuration

This section provides the configuration steps to configure the example illustrated in [Figure 14-1](#).

## Step 1 Enable IVR.

```
mds# conf t
Enter configuration commands, one per line. End with CNTL/Z.
mds (config)# ivr enable
mds (config)# exit
```

## Step 2 Verify that IVR Is enabled.

```
mds# show ivr
Inter-VSAN Routing is enabled

Inter-VSAN enabled switches

No IVR-enabled VSAN is active. Check VSAN-Topology configuration.

Inter-VSAN topology status

Current Status: Inter-VSAN topology is INACTIVE

Inter-VSAN zoneset status

 name :
 state : idle
 last activate time :
```

## Step 3 Configure the IVR VSAN-topology. In [Figure 14-1](#), two of the four IVR-enabled switches are members of VSANs 1 and 4. The other two switches are members of VSANs 2, 3, and 4.

```
mds# conf t
Enter configuration commands, one per line. End with CNTL/Z.
mds (config)# ivr vsan-topology database
mds (config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:05:40:01:1b:c2
vsan-ranges 1,4
mds (config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:02:00:44:22:00:4a:08
vsan-ranges 1,4
mds (config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:44:22:02:8a:04
vsan-ranges 2-4
mds (config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:44:22:40:aa:16
vsan-ranges 2-4
mds (config-ivr-topology-db)# exit
```

## Step 4 Verify the configured VSAN-topology.



### Note

The configured topology has not yet been activated—as indicated by the `no` status displayed in the Active column.

```
mds (config)# do show ivr vsan-topology
```

| AFID | SWITCH WWN                | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1    | 20:00:00:05:40:01:1b:c2 * | no     | yes 1,4    |
| 1    | 20:00:00:44:22:00:4a:08   | no     | yes 1,4    |
| 1    | 20:00:00:44:22:02:8a:04   | no     | yes 2-4    |
| 1    | 20:00:00:44:22:40:aa:16   | no     | yes 2-4    |



Total: 4 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is INACTIVE

**Step 5** Activate the Configured VSAN topology.

```
mds(config)# ivr vsan-topology activate
```

**Step 6** Verify the Activation.

```
mds(config)# do show ivr vsan-topology
```

| AFID | SWITCH                    | WWN | Active | Cfg. VSANS |
|------|---------------------------|-----|--------|------------|
| 1    | 20:00:00:05:40:01:1b:c2 * | yes | yes    | 1,4        |
| 1    | 20:00:00:44:22:00:4a:08   | yes | yes    | 1,4        |
| 1    | 20:00:00:44:22:02:8a:04   | yes | yes    | 2-4        |
| 1    | 20:00:00:44:22:40:aa:16   | yes | yes    | 2-4        |

Total: 4 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE

Last activation time: Tue May 20 23:14:59 1980

**Step 7** Configure IVR zoneset and zones. Two zones are required:

- one zone with the Tape T (pwwn 10:02:50:45:32:20:7a:52) and Server S1 (pwwn 10:02:66:45:00:20:89:04)
- another zone with the Tape T and Server S2 (pwwn 10:00:ad:51:78:33:f9:86)



**Tip**

Instead of creating two IVR zones, you can also create one IVR zone with the tape and both servers.

```
mds(config)# ivr zoneset name tape_server1_server2
```

```
mds(config-ivr-zoneset)# zone name tape_server1
```

```
mds(config-ivr-zoneset-zone)# member pwwn 10:02:50:45:32:20:7a:52 vsan 1
```

```
mds(config-ivr-zoneset-zone)# member pwwn 10:02:66:45:00:20:89:04 vsan 2
```

```
mds(config-ivr-zoneset-zone)# exit
```

```
mds(config-ivr-zoneset)# zone name tape_server2
```

```
mds(config-ivr-zoneset-zone)# member pwwn 10:02:50:45:32:20:7a:52 vsan 1
```

```
mds(config-ivr-zoneset-zone)# member pwwn 10:00:ad:51:78:33:f9:86 vsan 3
```

```
mds(config-ivr-zoneset-zone)# exit
```

**Step 8** View the IVR zone configuration to confirm that the IVR zoneset and IVR zones are properly configured.

```
mds(config)# do show ivr zoneset
```

```
zoneset name tape_server1_server2
```

```
zone name tape_server1
```

```
pwwn 10:02:50:45:32:20:7a:52 vsan 1
```

```
pwwn 10:02:66:45:00:20:89:04 vsan 2
```

```
zone name tape_server2
```

```
pwwn 10:02:50:45:32:20:7a:52 vsan 1
```

```
pwwn 10:00:ad:51:78:33:f9:86 vsan 3
```

**Step 9** View the zoneset prior to IVR zoneset activation. Prior to activating the IVR zoneset, view the active zoneset. Repeat this step for VSANs 2 and 3.

```
mds(config)# do show zoneset active vsan 1
```

```
zoneset name finance_dept vsan 1
```

```
zone name accounts_database vsan 1
```

```

pwwn 10:00:23:11:ed:f6:23:12
pwwn 10:00:56:43:11:56:fe:ee

zone name $default_zone$ vsan 1

```

**Step 10** Activate the configured IVR zoneset.

```

mds(config)# ivr zoneset activate name tape_server1_server2
zoneset activation initiated. check inter-VSAN zoneset status
mds(config)# exit

```

**Step 11** Verify the IVR zoneset activation.

```

mds# show ivr zoneset active
zoneset name tape_server1_server2
 zone name tape_server1
 pwwn 10:02:50:45:32:20:7a:52 vsan 1
 pwwn 10:02:66:45:00:20:89:04 vsan 2

 zone name tape_server2
 pwwn 10:02:50:45:32:20:7a:52 vsan 1
 pwwn 10:00:ad:51:78:33:f9:86 vsan 3

```

**Step 12** Verify the zoneset updates. Upon successful IVR zoneset activation, verify that appropriate zones are added to the active zoneset. Repeat this step for VSANs 2 and 3.

```

mds# show zoneset active vsan 1
zoneset name finance_dept vsan 1
 zone name accounts_database vsan 1
 pwwn 10:00:23:11:ed:f6:23:12
 pwwn 10:00:56:43:11:56:fe:ee

 zone name IVRZ_tape_server1 vsan 1
 pwwn 10:02:66:45:00:20:89:04
 pwwn 10:02:50:45:32:20:7a:52

 zone name IVRZ_tape_server2 vsan 1
 pwwn 10:02:50:45:32:20:7a:52
 pwwn 10:00:ad:51:78:33:f9:86

 zone name $default_zone$ vsan 1

mds# show ivr zoneset status
Zoneset Status

 name : tape_server1_server2
 state : activation success
 last activate time : Tue May 20 23:23:01 1980
 force option : on

status per vsan:

 vsan status
 ---- -
 1 active

```



# CHAPTER 15

## Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login database, the name server features, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family. It includes the following sections:

- [Displaying FLOGI Details, page 15-1](#)
- [Configuring the Name Server Proxy Feature, page 15-3](#)
- [Displaying FDMI, page 15-6](#)
- [Displaying RSCN Information, page 15-8](#)

## Displaying FLOGI Details

In a Fibre Channel fabric, each host or disk requires an FC ID. Use the **show flogi** command to verify if a storage device is displayed in the Fabric login (FLOGI) table as in the following examples. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports. See Examples 15-1 to 15-4.

### Example 15-1 Displays Details on the FLOGI Database

```
switch# show flogi database
```

| INTERFACE | VSAN | FCID     | PORT NAME               | NODE NAME               |
|-----------|------|----------|-------------------------|-------------------------|
| sup-fc0   | 2    | 0xb30100 | 10:00:00:05:30:00:49:63 | 20:00:00:05:30:00:49:5e |
| fc9/13    | 1    | 0xb200e2 | 21:00:00:04:cf:27:25:2c | 20:00:00:04:cf:27:25:2c |
| fc9/13    | 1    | 0xb200e1 | 21:00:00:04:cf:4c:18:61 | 20:00:00:04:cf:4c:18:61 |
| fc9/13    | 1    | 0xb200d1 | 21:00:00:04:cf:4c:18:64 | 20:00:00:04:cf:4c:18:64 |
| fc9/13    | 1    | 0xb200ce | 21:00:00:04:cf:4c:16:fb | 20:00:00:04:cf:4c:16:fb |
| fc9/13    | 1    | 0xb200cd | 21:00:00:04:cf:4c:18:f7 | 20:00:00:04:cf:4c:18:f7 |

Total number of flogi = 6.

### Example 15-2 Displays the FLOGI Database by Interface

```
switch# show flogi database interface fc1/11
```

| INTERFACE | VSAN | FCID     | PORT NAME               | NODE NAME               |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/11    | 1    | 0xa002ef | 21:00:00:20:37:18:17:d2 | 20:00:00:20:37:18:17:d2 |

## ■ Displaying FLOGI Details

|        |   |          |                         |                         |
|--------|---|----------|-------------------------|-------------------------|
| fc1/11 | 1 | 0xa002e8 | 21:00:00:20:37:38:a7:c1 | 20:00:00:20:37:38:a7:c1 |
| fc1/11 | 1 | 0xa002e4 | 21:00:00:20:37:6b:d7:18 | 20:00:00:20:37:6b:d7:18 |
| fc1/11 | 1 | 0xa002e2 | 21:00:00:20:37:18:d2:45 | 20:00:00:20:37:18:d2:45 |
| fc1/11 | 1 | 0xa002e1 | 21:00:00:20:37:39:90:6a | 20:00:00:20:37:39:90:6a |
| fc1/11 | 1 | 0xa002e0 | 21:00:00:20:37:36:0b:4d | 20:00:00:20:37:36:0b:4d |
| fc1/11 | 1 | 0xa002dc | 21:00:00:20:37:5a:5b:27 | 20:00:00:20:37:5a:5b:27 |
| fc1/11 | 1 | 0xa002da | 21:00:00:20:37:18:6f:90 | 20:00:00:20:37:18:6f:90 |
| fc1/11 | 1 | 0xa002d9 | 21:00:00:20:37:5b:cf:b9 | 20:00:00:20:37:5b:cf:b9 |
| fc1/11 | 1 | 0xa002d6 | 21:00:00:20:37:46:78:97 | 20:00:00:20:37:46:78:97 |

Total number of flogi = 10.

**Example 15-3 Displays the FLOGI Database by VSAN**

```
switch# show flogi database vsan 1
```

| INTERFACE | VSAN | FCID     | PORT NAME               | NODE NAME               |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/3     | 1    | 0xef02ef | 22:00:00:20:37:18:17:d2 | 20:00:00:20:37:18:17:d2 |
| fc1/3     | 1    | 0xef02e8 | 22:00:00:20:37:38:a7:c1 | 20:00:00:20:37:38:a7:c1 |
| fc1/3     | 1    | 0xef02e4 | 22:00:00:20:37:6b:d7:18 | 20:00:00:20:37:6b:d7:18 |
| fc1/3     | 1    | 0xef02e2 | 22:00:00:20:37:18:d2:45 | 20:00:00:20:37:18:d2:45 |
| fc1/3     | 1    | 0xef02e1 | 22:00:00:20:37:39:90:6a | 20:00:00:20:37:39:90:6a |
| fc1/3     | 1    | 0xef02e0 | 22:00:00:20:37:36:0b:4d | 20:00:00:20:37:36:0b:4d |
| fc1/3     | 1    | 0xef02dc | 22:00:00:20:37:5a:5b:27 | 20:00:00:20:37:5a:5b:27 |
| fc1/3     | 1    | 0xef02da | 22:00:00:20:37:18:6f:90 | 20:00:00:20:37:18:6f:90 |
| fc1/3     | 1    | 0xef02d9 | 22:00:00:20:37:5b:cf:b9 | 20:00:00:20:37:5b:cf:b9 |
| fc1/3     | 1    | 0xef02d6 | 22:00:00:20:37:46:78:97 | 20:00:00:20:37:46:78:97 |

Total number of flogi = 10.

**Example 15-4 Displays the FLOGI Database by FC ID**

```
switch# show flogi database fcid 0xef02e2
```

| INTERFACE | VSAN | FCID     | PORT NAME               | NODE NAME               |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/3     | 1    | 0xef02e2 | 22:00:00:20:37:18:d2:45 | 20:00:00:20:37:18:d2:45 |

Total number of flogi = 1.

See the [“Allocating Flat FC IDs”](#) section on page 29-19 and the [“Enabling Loop Monitoring”](#) section on page 29-20.

## Configuring the Name Server Proxy Feature

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device which originally registered the information.

The proxy feature is useful when you wish to modify (update or delete) the contents of a database entry that was previously registered by a different device.

All name server registration requests come from the same port whose parameter is registered or changed. If it doesn't, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

## Registering Name Server Proxies

To register the name server proxy, follow these steps:

|        | Command                                                                                  | Purpose                                         |
|--------|------------------------------------------------------------------------------------------|-------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                               | Enters configuration mode.                      |
| Step 2 | switch(config)# <b>fcns proxy-port 21:00:00:e0:8b:00:26:d0 vsan 2</b><br>switch(config)# | Configures a proxy port for the specified VSAN. |

## Displaying Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

Use the **show fcns** command to display the name server database and statistical information for a specified VSAN or for all VSANs (see Examples 15-5 to 15-8).

### Example 15-5 Displays the Name Server Database

```
switch# show fcns database
```

```

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x010000 N 50:06:0b:00:00:10:a7:80 (Andiamo) scsi-fcp fc-gs
0x010001 N 10:00:00:05:30:00:24:63 (Andiamo) ipfc
0x010002 N 50:06:04:82:c3:a0:98:52 (Company 1) scsi-fcp 250
0x010100 N 21:00:00:e0:8b:02:99:36 (Company A) scsi-fcp
0x020000 N 21:00:00:e0:8b:08:4b:20 (Company A)
0x020100 N 10:00:00:05:30:00:24:23 (Andiamo) ipfc
0x020200 N 21:01:00:e0:8b:22:99:36 (Company A) scsi-fcp
```

**Example 15-6 Displays the Name Server Database for the Specified VSAN**

```
switch# show fcns database vsan 1
VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x030001 N 10:00:00:05:30:00:25:a3 (Cisco) ipfc
0x030101 NL 10:00:00:00:77:99:60:2c (Interphase)
0x030200 N 10:00:00:49:c9:28:c7:01
0x030300 N 10:00:00:4a:c9:28:c7:01
0x030400 N 10:00:00:59:c9:28:c7:01
0xec0001 NL 21:00:00:20:37:a6:be:14 (Seagate) scsi-fcp
0xec0100 N 10:00:00:05:30:00:26:23 (Cisco) ipfc
0xec0200 N 10:00:00:5a:c9:28:c7:01

Total number of entries = 8
```

**Example 15-7 Displays the Name Server Database Details**

```
switch# show fcns database detail

VSAN:1 FCID:0x030001

port-wwn (vendor) :10:00:00:05:30:00:25:a3 (Cisco)
node-wwn :20:00:00:05:30:00:25:9e
class :2,3
node-ip-addr :0.0.0.0
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:ipfc
symbolic-port-name :
symbolic-node-name :
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :00:00:00:00:00:00:00:00
hard-addr :0x000000
.
.
.

VSAN:1 FCID:0xec0200

port-wwn (vendor) :10:00:00:5a:c9:28:c7:01
node-wwn :10:00:00:5a:c9:28:c7:01
class :3
node-ip-addr :0.0.0.0
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:
symbolic-port-name :
symbolic-node-name :
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :22:0a:00:05:30:00:26:1e
hard-addr :0x000000

Total number of entries = 8
```

**Example 15-8 Displays the Name Server Statistics**

```
switch# show fcns statistics
registration requests received = 27
deregistration requests received = 0
queries received = 57
queries sent = 10
reject responses sent = 14
RSCNs received = 0
RSCNs sent = 0
```

# Displaying FDMI

SAN-OS 1.3(x) provides support for the Fabric-Device Management Interface (FDMI) functionally, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel Host Bus Adaptors (HBAs) through inband communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the SAN-OS software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

Use the **show fdmi** command to display the FDMI database information (see Examples 15-9 to 15-11).

## Example 15-9 Displays All HBA Management Servers

```
switch# show fdmi database
Registered HBA List for VSAN 1
 10:00:00:00:c9:32:8d:77
 21:01:00:e0:8b:2a:f6:54
switch# show fdmi database detail
Registered HBA List for VSAN 1

HBA-ID: 10:00:00:00:c9:32:8d:77

Node Name :20:00:00:00:c9:32:8d:77
Manufacturer :Emulex Corporation
Serial Num :0000c9328d77
Model :LP9002
Model Description :Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver :2002606D
Driver Ver :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver :3.11A0
Firmware Ver :3.90A7
OS Name/Ver :Window 2000
CT Payload Len :1300000
 Port-id: 10:00:00:00:c9:32:8d:77

HBA-ID: 21:01:00:e0:8b:2a:f6:54

Node Name :20:01:00:e0:8b:2a:f6:54
Manufacturer :QLogic Corporation
Serial Num :74262
Model :QLA2342
Model Description :QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver :FC5010409-10
Driver Ver :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver :1.24
Firmware Ver :03.02.13.
OS Name/Ver :500
CT Payload Len :2040
 Port-id: 21:01:00:e0:8b:2a:f6:54
```



**Example 15-10 Displays HBA Details for a Specified VSAN**

```

switch# show fdbi database detail vsan 1
Registered HBA List for VSAN 1

HBA-ID: 10:00:00:00:c9:32:8d:77

Node Name :20:00:00:00:c9:32:8d:77
Manufacturer :Emulex Corporation
Serial Num :0000c9328d77
Model :LP9002
Model Description:Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver :2002606D
Driver Ver :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver :3.11A0
Firmware Ver :3.90A7
OS Name/Ver :Window 2000
CT Payload Len :1300000
Port-id: 10:00:00:00:c9:32:8d:77

HBA-ID: 21:01:00:e0:8b:2a:f6:54

Node Name :20:01:00:e0:8b:2a:f6:54
Manufacturer :QLogic Corporation
Serial Num :\74262
Model :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver :FC5010409-10
Driver Ver :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver :1.24
Firmware Ver :03.02.13.
OS Name/Ver :500
CT Payload Len :2040
Port-id: 21:01:00:e0:8b:2a:f6:54

```

**Example 15-11 Displays Details for the Specified HBA Entry**

```

switch# show fdbi database detail Hba-id 21:01:00:e0:8b:2a:f6:54 vsan 1

Node Name :20:01:00:e0:8b:2a:f6:54
Manufacturer :QLogic Corporation
Serial Num :\74262
Model :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver :FC5010409-10
Driver Ver :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver :1.24
Firmware Ver :03.02.13.
OS Name/Ver :500
CT Payload Len :2040
Port-id: 21:01:00:e0:8b:2a:f6:54

```

# Displaying RSCN Information

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through SCR). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric.
- A name server registration change.
- A new zone enforcement.
- IP address change
- Or any other similar event that affects the operation of the host

Apart from sending these events to registered hosts a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.



## Note

The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the Name Server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

Use the **show rscn** command to display RSCN information (see Examples 15-12 and 15-13).

### Example 15-12 Displays Register Device Information

```
switch# show rscn scr-table vsan 1
SCR table for VSAN: 1

FC-ID REGISTERED FOR

0x1b0300 fabric detected rscns

Total number of entries = 1
```



## Note

The SCR table cannot be configured, it is only populated if one or more hosts send SCR frames to register for RSCN information. If the **show rscn scr-table** command does not return any entries, no host is interested in receiving RSCN information.

### Example 15-13 Displays RSCN Counter Information

```
switch# show rscn statistics vsan 1

Statistics for VSAN: 1

Number of SCR received = 8
Number of SCR ACC sent = 8
Number of SCR RJT sent = 0
Number of RSCN received = 0
Number of RSCN sent = 24
Number of RSCN ACC received = 24
Number of RSCN ACC sent = 0
Number of RSCN RJT received = 0
Number of RSCN RJT sent = 0
Number of SW-RSCN received = 6
Number of SW-RSCN sent = 15
```

```

Number of SW-RSCN ACC received = 15
Number of SW-RSCN ACC sent = 6
Number of SW-RSCN RJT received = 0
Number of SW-RSCN RJT sent = 0

```

## Sending RSCNs

If the RSCN **multi-pid** option is enabled then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs.

For example, you have two disks (D1, D2) and a host (Host H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2 and H belong to the same zone. If disks D1 and D2 are online at the same time, then one of the following applies:

- If the **multi-pid** option is disabled on switch 1, then two RSCNs is generated to Host H—one for the disk D1 and another for disk D2.
- If the **multi-pid** format is enabled on switch 1, then a single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).

To configure the **multi-pid** option, follow these steps:

|        | Command                                        | Purpose                                                   |
|--------|------------------------------------------------|-----------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#     | Enters configuration mode.                                |
| Step 2 | switch(config)# <b>rscn multi-pid vsan 105</b> | Sends RSCNs in a multi-pid format for the specified VSAN. |



### Note

Some Nx ports may not understand multi-pid RSCN payloads. If so, you must disable the **multi-pid** RSCN option.

## Clearing RSCN Statistics

Use the **clear rscn statistics** command to clear the RSCN statistics for the specified VSAN.

```
switch# clear rscn statistics vsan 1
```

After clearing the RSCN statistics, you can view the cleared counters by issuing the **show rscn** command.

```
switch# show rscn statistics vsan 1
```

```

Statistics for VSAN: 1

```

```

Number of SCR received = 0
Number of SCR ACC sent = 0
Number of SCR RJT sent = 0
Number of RSCN received = 0
Number of RSCN sent = 0
Number of RSCN ACC received = 0
Number of RSCN ACC sent = 0
Number of RSCN RJT received = 0
Number of RSCN RJT sent = 0

```

```
Number of SW-RSCN received = 0
Number of SW-RSCN sent = 0
Number of SW-RSCN ACC received = 0
Number of SW-RSCN ACC sent = 0
Number of SW-RSCN RJT received = 0
Number of SW-RSCN RJT sent = 0
```

This command is used for debugging purposes. When you clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (like ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.



## Configuring Switch Security

---

The authentication, authorization, and accounting (AAA) strategy is used to verify identity of, grant access, and track the actions of remote users in all switches in the Cisco MDS 9000 Family. The Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) provide AAA solutions.

Based on the user ID and password combination provided, switches perform local authentication using a local database or remote authentication using AAA server(s). A global, preshared, secret key authenticates communication between the AAA servers. This secret key can be configured for all AAA server groups or for only a specific AAA server. This kind of authentication provides a central configuration management capability.

This chapter includes the following sections:

- [Switch Management Security, page 16-2](#)
- [Switch AAA Functionalities, page 16-3](#)
- [Configuring RADIUS, page 16-6](#)
- [Configuring TACACS+, page 16-10](#)
- [Configuring Server Groups, page 16-14](#)
- [Local AAA, page 16-15](#)
- [No AAA Authentication, page 16-15](#)
- [Displaying AAA Authentication, page 16-15](#)
- [Authentication and Authorization Process, page 16-16](#)
- [Configuring Role-Based CLI Authorization, page 16-18](#)
- [Configuring CLI User Profiles, page 16-22](#)
- [Configuring CLI Accounting Parameters, page 16-24](#)
- [Recovering Administrator Password, page 16-26](#)
- [Configuring SSH Services, page 16-27](#)
- [SNMP Security, page 16-30](#)
- [Default Security Settings, page 16-36](#)

# Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family is implemented using the Command-line interface (CLI) or Simple Network Management Protocol (SNMP):

## CLI

You can access the CLI using the Console (serial connection), Telnet, or Secure Shell (SSH). For each management path (console or Telnet and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
  - Using Remote Authentication Dial-In User Services (RADIUS). See the [“Configuring RADIUS” section on page 16-6](#).
  - Using Terminal Access Controller Access Control System plus (TACACS+). See the [“Configuring TACACS+” section on page 16-10](#)
- Local security control. See the [“Local AAA” section on page 16-15](#)
- Trivial authentication. See the [“No AAA Authentication” section on page 16-15](#)

These authentication mechanisms can also be used to configure AAA for the following scenarios:

- iSCSI authentication (see the [Authentication Mechanism, page 22-61](#) section in the [Chapter 22, “Configuring IP Storage.”](#))
- Fibre Channel Security Protocol (FC-SP) authentication (see the [Chapter 17, “Configuring Fabric Security.”](#))

## SNMP Security

The SNMP agent supports security features for SNMPv1, SNMPv 2c, and SNMPv3. Normal SNMP security mechanisms apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).

See the [“SNMP Security” section on page 16-30](#).

Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for information on the Fabric Manager.



### Note

Users and roles configured through the CLI are different from users and roles configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

# Switch AAA Functionalities

Using CLI, you can configure Authentication, Authorization, and Accounting (AAA) switch functionalities on any switch in the Cisco MDS 9000 Family. This section includes the following topics:

- [Authentication, page 16-3](#)
- [Authorization, page 16-3](#)
- [Accounting, page 16-4](#)
- [Remote Authentication by AAA Servers, page 16-4](#)
- [Remote Authentication Guidelines, page 16-4](#)
- [Server Groups, page 16-4](#)

## Authentication

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

- See the [“Configuring RADIUS” section on page 16-6](#) for more information on configuring RADIUS authentication
- See the [“Configuring TACACS+” section on page 16-10](#) section for more information on configuring TACACS+ authentication.
- See the [“Local AAA” section on page 16-15](#) section for more information on configuring local AAA authentication.
- See the [“No AAA Authentication” section on page 16-15](#) section for more information on not configuring AAA authentication.

## Authorization

By default, two roles exist in all switches:

- Network operator (`network-operator`)—Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator (`network-admin`)—Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.

The two default roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Assign user roles either locally or using remote AAA servers (see the [“Configuring Role-Based CLI Authorization” section on page 16-18](#)).
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when that user is authenticated through remote AAA server (see the [“Configuring CLI User Profiles” section on page 16-22](#)).

## Accounting

Accounting refers to the log that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally and remotely. See the [“Configuring CLI Accounting Parameters” section on page 16-24](#) for log size information.

## Remote Authentication by AAA Servers

AAA authentication provides the following advantages over local database authentication:

- Requires only one password to be shared between the switch and the AAA servers.
- Easier to manage user password lists for each switch in the fabric.
- AAA servers are deployed widely across enterprises and can be easily adopted.

## Remote Authentication Guidelines

When you prefer using remote C servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- If all configured AAA servers are not reachable, the policy configured on the switch determines the authentication method.
- RADIUS servers are easily reachable if an overlay Ethernet LAN is attached to the switch (see [Chapter 22, “Configuring IP Storage”](#)). This is the recommended method.
- SAN networks connected to the switch should have at least one gateway switch connect to the Ethernet LAN containing the AAA servers

**Note**

---

If you are using IP connectivity to reach an AAA server, the SAN connects to the switch.

---

## Server Groups

You can specify remote AAA servers for authentication, authorization and accounting using server groups. A server group consists of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for fail-over servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fails to respond, then that server group option is considered a failure. You can create a server group using the **aaa group server** command.

If required, you can specify multiple server groups. If the MDS switch encounters errors from the server(s) in the first group, it tries the servers in next server group.



## AAA Service Configuration Options

AAA configuration in Cisco MDS switches is service based. You can have separate AAA configurations for following services

- Telnet or SSH login—use the **aaa authentication login default** command
- Console login—use the **aaa authentication login console** command
- iSCSI authentication—use the **aaa authentication iscsi default** command (see the [Authentication Mechanism, page 22-61](#) section in the [Chapter 22, “Configuring IP Storage.”](#))
- FC-SP authentication—use the **aaa authentication dhchap default** command (see the [Chapter 17, “Configuring Fabric Security.”](#))
- Accounting—use the **aaa accounting default** command

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option will be tried in the order specified. If all the methods fail, local is tried

**Note**

---

Even if local is not specified as one of the options, it is tried when all other configured options fail.

---

# Configuring RADIUS

Cisco MDS switches use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and set timeout and retry counts.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities. This section includes the following topics:

- [About RADIUS, page 16-6](#)
- [Setting the RADIUS Server Address, page 16-6](#)
- [Setting the RADIUS Preshared Key, page 16-7](#)
- [Setting the RADIUS Server Time-Out Interval, page 16-8](#)
- [Setting Iterations of the RADIUS Server, page 16-8](#)
- [Defining Vendor-Specific Attributes, page 16-8](#)
- [Displaying RADIUS Server Details, page 16-9](#)

## About RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

RADIUS is a fully open protocol, distributed in source code format, that can be modified to work with any security system currently available on the market.

You can set the RADIUS server address, the RADIUS preshared key, the RADIUS server time-out interval, iterations of the RADIUS server, define vendor-specific attributes, and display RADIUS server details.

## Setting the RADIUS Server Address

You can add up to 64 RADIUS servers using the **radius-server host** command. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys.

To specify the RADIUS server address and the options, follow these steps:

|        | Command                                                                                         | Purpose                                                                                                                                                                                                                                                       |
|--------|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>conf t</b>                                                                           | Enters configuration mode.                                                                                                                                                                                                                                    |
| Step 2 | switch(config)# <b>radius-server host 10.10.0.0</b><br><b>key HostKey</b><br>switch(config)#    | Specifies a key for the selected RADIUS server. This key overrides the key assigned using the <b>radius-server key</b> command. In this example, the host is 10.10.0.0 and the key is HostKey.                                                                |
| Step 3 | switch(config)# <b>radius-server host 10.10.0.0</b><br><b>auth-port 2003</b><br>switch(config)# | Specifies the destination UDP port number to which the RADIUS authentication messages should be sent. In this example, the host is 10.10.0.0 and the authentication port is 2003. The default authentication port is 1812, and the valid range is 0 to 65366. |

|        | Command                                                                                   | Purpose                                                                                                                                                                                                                                            |
|--------|-------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 4 | switch(config)# <b>radius-server host 10.10.0.0<br/>acct-port 2004</b><br>switch(config)# | Specifies the destination UDP port number to which RADIUS accounting messages should be sent. The default accounting port is 1813, and the valid range is 0 to 65366.                                                                              |
| Step 5 | switch(config)# <b>radius-server host 10.10.0.0<br/>accounting</b><br>switch(config)#     | Specifies this server to be used only for accounting purposes.<br><br><b>Note</b> If neither the <b>authentication</b> option nor the <b>accounting</b> options are specified, the server is used for both accounting and authentication purposes. |
| Step 6 | switch(config)# <b>radius-server host radius2<br/>key 0 abcd</b><br>switch(config)#       | Specifies a clear text key for the specified server. The key is restricted to 65 characters.                                                                                                                                                       |
|        | switch(config)# <b>radius-server host radius3<br/>key 7 1234</b><br>switch(config)#       | Specifies a reversible encrypted key for the specified server. The key is restricted to 65 characters.                                                                                                                                             |

## Setting the RADIUS Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 65 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch. You can override this global key assignment by explicitly using the **key** option in the **radius-server host** command.

To set the RADIUS preshared key, follow these steps:

|        | Command                                                                   | Purpose                                                                                                                                              |
|--------|---------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                   | Enters configuration mode.                                                                                                                           |
| Step 2 | switch(config)# <b>radius-server key AnyWord</b><br>switch(config)#       | Configures a preshared key (AnyWord) to authenticate communication between the RADIUS client and server. The default is clear text.                  |
|        | switch(config)# <b>radius-server key 0<br/>AnyWord</b><br>switch(config)# | Configures a preshared key (AnyWord) specified in clear text (indicated by 0) to authenticate communication between the RADIUS client and server.    |
|        | switch(config)# <b>radius-server key 7 public</b><br>switch(config)#      | Configures a preshared key (public) specified in encrypted text (indicated by 7) to authenticate communication between the RADIUS client and server. |

## Setting the RADIUS Server Time-Out Interval

To specify the time between retransmissions to the RADIUS servers, follow these steps:

|        | Command                                                                               | Purpose                                                                                                                                                     |
|--------|---------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                         | Enters configuration mode.                                                                                                                                  |
| Step 2 | <code>switch(config)# radius-server timeout 30</code><br><code>switch(config)#</code> | Specifies the time (in seconds) between retransmissions to the RADIUS server. The default time-out is one (1) second. The time range in seconds is 1 to 60. |

You can revert the retransmission time to its default by issuing the **no radius-server timeout** command.

## Setting Iterations of the RADIUS Server

By default, a switch retries a RADIUS server connection only once. This number can be configured. The maximum is five retries per server. You can revert the retry number to its default by issuing the **no radius-server retransmit** command.

To specify the number of times that RADIUS servers should try to authenticate a user, follow these steps:

|        | Command                                                                                 | Purpose                                                                                                                        |
|--------|-----------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                           | Enters configuration mode.                                                                                                     |
| Step 2 | <code>switch(config)# radius-server retransmit 3</code><br><code>switch(config)#</code> | Configures the number of times (3) the switch tries to connect to a RADIUS server(s) before reverting to local authentication. |

## Defining Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named `cisco-avpair`. The value is a string with the following format:

```
protocol : attribute sep value *
```

Where `protocol` is a Cisco attribute for a particular type of authorization, and `sep` is = for mandatory attributes, and \* is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, like authorization information, along with authentication results. This authorization information is specified through VSAs.

### VSA Format

The following VSA protocol options are supported:

- Shell protocol—used in Access-Accept packets to provide user profile information.

- Accounting protocol—used in Accounting-Request packets. If a value contains any white spaces, it should be put within double quotation marks.

The following attributes are supported:

- `roles`—This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles `vsan-admin` and `storage-admin`, the value field would be “`vsan-admin storage-admin`.” This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. This is an example using the roles attribute:

```
Cisco-AVPair = shell:roles="network-admin vsan-admin"
```

- `accountinginfo`—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol value.

## Displaying RADIUS Server Details

Use the **show radius-server** command to display configured RADIUS parameters (see [Example 16-1](#)).



### Note

Only administrators can view the RADIUS preshared key.

### Example 16-1 Displays Configured RADIUS Information

```
switch# show radius-server
Global RADIUS shared secret:Myxgqc
retransmission count:5
timeout value:10
following RADIUS servers are configured:
 myradius.cisco.users.com:
 available for authentication on port:1812
 available for accounting on port:1813
 172.22.91.37:
 available for authentication on port:1812
 available for accounting on port:1813
 RADIUS shared secret:23MHcUnD
 10.10.0.0:
 available for authentication on port:1812
 available for accounting on port:1813
 RADIUS shared secret:hostkey----> for administrators only
```

### Example 16-2 Displays Configured RADIUS Server-Group Order

```
switch# show radius-server groups
total number of groups:4
following RADIUS server groups are configured:
 group radius:
 server: all configured radius servers
 group Group1:
 server: Server3 on auth-port 1812, acct-port 1813
 server: Server5 on auth-port 1812, acct-port 1813
 group Group5:
```

# Configuring TACACS+

A Cisco MDS switch uses the Terminal Access Controller Access Control System plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values. This section includes the following topics:

- [About TACACS+, page 16-10](#)
- [Advantages of TACACS+, page 16-10](#)
- [Enabling TACACS+, page 16-11](#)
- [Setting the TACACS+ Server Address, page 16-11](#)
- [Setting the Secret Key, page 16-12](#)
- [Setting the Timeout Value, page 16-12](#)
- [Defining Custom Attributes for Roles, page 16-12](#)
- [Displaying TACACS+ Server Details, page 16-13](#)

## About TACACS+

TACACS+ is a client-server protocol which uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The addition of TACACS+ support in SAN-OS 1.3(x) enables the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities—authorization can be done without authentication.
- Performs independent of servers if it is configured to its own database.
- TCP transport protocol to send data between the AAA client and server, using reliable transfers with a connection-oriented protocol
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality—the RADIUS protocol only encrypts passwords.

## Advantages of TACACS+

This section provides a brief list of advantages that TACACS+ has over and RADIUS.

- Uses TCP protocol which has a connection-oriented transport
- Provides built-in transport support
- Provides a separate acknowledgment that a request has been received
- Provides immediate indication of a crashed, or not running, server
- Detects server crashes out-of-band with actual requests
- Maintains simultaneous connections to multiple servers
- Adapts to growing, as well as congested networks

For a detailed comparison visit the following URL:

<http://www.cisco.com/warp/public/480/10.html#comparing>

## Enabling TACACS+

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the TACACS+ feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

To enable TACACS+ for a Cisco MDS switch, follow these steps:

|        | Command                                  | Purpose                                        |
|--------|------------------------------------------|------------------------------------------------|
| Step 1 | switch# <b>config t</b>                  | Enters configuration mode.                     |
| Step 2 | switch(config)# <b>tacacs+ enable</b>    | Enables the TACACS+ in this switch.            |
|        | switch(config)# <b>no tacacs+ enable</b> | Disables (default) the TACACS+ in this switch. |

## Setting the TACACS+ Server Address

Use the **tacacs-server** command to configure the communication parameters for the required TACACS+ server.

If a secret key is not configured for a configured server, a warning message is issued and the global secret encryption key is automatically used (see the [“Setting the Secret Key”](#) section on page 16-12).

To configure TACACS+ server option, follow these steps:

|        | Command                                                                                              | Purpose                                                                                                                             |
|--------|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                              | Enters configuration mode.                                                                                                          |
| Step 2 | switch(config)# <b>tacacs-server host 171.71.58.91</b><br>warning: no key is configured for the host | Configures the TACACS+ server identified by the specified IP address.                                                               |
|        | switch(config)# <b>no tacacs-server host 10.10.1.0</b>                                               | Deletes the specified TACACS+ server identified by the IP address. By default, no server is configured.                             |
| Step 3 | switch(config)# <b>tacacs-server host 171.71.58.91 port 2</b>                                        | Configures the TCP port for all TACACS+ requests.                                                                                   |
|        | switch(config)# <b>no tacacs-server host 171.71.58.91 port 2</b>                                     | Reverts to the factory default of using Port 49 for server access.                                                                  |
| Step 4 | switch(config)# <b>tacacs-server host host1.cisco.com key MyKey</b>                                  | Configures the TACACS+ server identified by the specified domain name and assigns a special key.                                    |
| Step 5 | switch(config)# <b>tacacs-server host host100.cisco.com timeout 25</b>                               | Configures the timeout period for the switch to wait for a response from the specified server before it declares a timeout failure. |

## Setting the Secret Key

Use the **tacacs-server** command to configure global values for the **key** for all TACACS+ servers.



### Note

Secret keys configured for individual servers override the globally configured values.

To set the secret key for TACACS+ servers, follow these steps:

|        | Command                                                          | Purpose                                                                                                                                                                                                                                                                                                                                                        |
|--------|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                          | Enters configuration mode.                                                                                                                                                                                                                                                                                                                                     |
| Step 2 | switch(config)# <b>tacacs-server key</b><br><b>7 tacacsPword</b> | Assigns the global secret key to access the TACACS+ server. This example specifies <b>7</b> to indicate encryption. If this global key and the individual server keys are not configured, clear text messages are sent to the TACACS+ server(s).<br><br><b>Note</b> If secret keys are configured for individual servers, those keys override this global key. |
|        | switch(config)# <b>no tacacs-server key oldPword</b>             | Deletes the configured secret key to access the TACACS+ server and reverts to the factory default of allowing access to all configured servers.                                                                                                                                                                                                                |

## Setting the Timeout Value

Use the **tacacs-server** command to configure global **timeout** values for all TACACS+ servers.



### Note

Timeout values configured for individual servers override the globally configured values.

To set the share password for TACACS+ servers, follow these steps:

|        | Command                                            | Purpose                                                                                                                                                                                                              |
|--------|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                            | Enters configuration mode.                                                                                                                                                                                           |
| Step 2 | switch(config)# <b>tacacs-server timeout 30</b>    | Configures the global timeout period for the switch to wait for a response from all servers before it declares a timeout failure.                                                                                    |
|        | switch(config)# <b>no tacacs-server timeout 30</b> | Deletes the configured timeout period and reverts to the factory default of 5 seconds.<br><br><b>Note</b> If the timeout value is configured for individual servers, that value overrides this global timeout value. |

## Defining Custom Attributes for Roles

MDS uses TACACS+ custom attribute for service shell to configure the roles to which a user belongs. TACACS+ attributes are specified as **name=value** format. The attribute name for this custom attribute is **cisco-av-pair**. The following example illustrates how to specify roles using this attribute:

```
cisco-av-pair=shell:roles="network-admin vsan-admin"
```



**Note**

TACACS+ custom attributes can be defined on an ACS server for various services (for example, shell). MDS requires the TACACS+ custom attribute for service shell to be used for defining roles.

## Displaying TACACS+ Server Details

Use the **show tacacs+** commands to display configurations for the TACACS+ protocol configuration in all switches in the Cisco MDS 9000 Family (see [Example 16-3](#)).

### *Example 16-3 Displays Configured TACACS+ Server Information*

```
switch# show tacacs-server
Global TACACS+ shared secret:tacacsPword
timeout value:30
total number of servers:3

following TACACS+ servers are configured:
 171.71.58.91:
 available on port:2
 cisco.com:
 available on port:49
 171.71.22.95:
 available on port:49
 TACACS+ shared secret:MyKey
```

### *Example 16-4 Displays Configured TACACS+ Server Groups*

```
switch# show tacacs-server groups
total number of groups:1

following TACACS+ server groups are configured:
 group TacServer:
 server 171.71.58.91 on port 2
```

### *Example 16-5 Displays AAA Authentication Information*

```
switch# show aaa authentication
default: group TacServer local none
console: local
iscsi: local
dhchap: local
```

### *Example 16-6 Displays Configured Server Groups*

```
switch# show tacacs-server groups
total number of groups:2

following TACACS+ server groups are configured:
 group TacServer:
 server 171.71.58.91 on port 2
 group TacacsServer1:
 server ServerA on port 49
 server ServerB on port 49:
```

**Example 16-7 Displays All AAA Server Groups**

```
switch# show aaa groups
radius
TacServer
```

## Configuring Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the same protocol: either RADIUS or TACACS+. The servers are tried in the same order in which you configure them.

You can configure these server groups at any time but they only take effect when you apply them to a AAA service using the **aaa authentication login** command (see the [Defining Custom Attributes for Roles](#), page 16-12).

You can specify one or more remote AAA servers to authenticate users using server groups.

To specify the TACACS+ server order within a group, follow these steps:

|        | Command                                                                                      | Purpose                                                                                                                                                                                                     |
|--------|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                      | Enters configuration mode.                                                                                                                                                                                  |
| Step 2 | switch(config)# <b>aaa group<br/>server tacacs+ TacacsServer1</b><br>switch(config-tacacs+)# | Configures a TacacsServer1 group and enters the submode for that group.                                                                                                                                     |
|        | switch(config)# <b>aaa group<br/>server tacacs+ TacacsServer19</b>                           | Deletes the group called TacacsServer1 from the authentication list.                                                                                                                                        |
| Step 3 | switch(config-tacacs+)#<br><b>server ServerA</b>                                             | Configures ServerA to be tried first within TacacsServer1.<br><br><b>Tip</b> If the specified TACACS+ server is not found, configure it using the <b>tacacs-server host</b> command and retry this command. |
| Step 4 | switch(config-radius)# <b>server<br/>ServerB</b>                                             | Configures ServerB to be tried second within TacacsServer1.                                                                                                                                                 |
|        | switch(config-radius)# <b>no<br/>server ServerZ</b>                                          | Deletes ServerZ within the TacacsServer1 list of servers.                                                                                                                                                   |

To verify the configured server-group order, use the **show tacacs-server groups** command:

```
switch# show tacacs-server groups
total number of groups:2

following TACACS+ server groups are configured:
 group TacServer:
 server 171.71.58.91 on port 2
 group TacacsServer1:
 server ServerA on port 49
 server ServerB on port 49:
```

# Local AAA

The system maintains the user name and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information. You can configure local users using the **username** command (see the “[Creating or Updating Users](#)” section on page 16-22), and view the local accounting log using the **show accounting log** command (see [Example 16-8](#)).

## Example 16-8 Displays the Accounting Log Information

```
switch# show accounting log

Sat Jan 24 03:22:06 1981:stop:snmp_349154526_171.71.58.69:admin:
Sat Jan 24 03:22:06 1981:start:snmp_349154526_171.71.58.69:admin:
Sat Jan 24 03:22:06 1981:update:snmp_349154526_171.71.58.69:admin:Added member [
 WWN: 21:00:00:20:37:a6:be:00 ID: 2] to zone test-27 on VSAN 1
...
Sat Jan 24 23:59:56 1981:stop:/dev/pts/0_349228792:root:shell terminated
Sun Jan 25 00:00:06 1981:start:/dev/pts/1_349228806:admin:
```

# No AAA Authentication

You can turn off password verification using the **none** option in the **aaa authentication login** command. If you configure this option, users will be able to login without giving a valid password. But the user should at least exist locally on the MDS switch.(created using the **username** command).



### Caution

Use this option cautiously. If configured, any user will be able to access the switch at any time.

# Displaying AAA Authentication

The **show aaa authentication** command displays the configured authentication methods (see [Example 16-8](#)).

## Example 16-9 Example 16-8 Displays Authentication Information

```
switch# show aaa authentication

No AAA Authentication
default: group TacServer local none
console: local none
iscsi: local
dhchap: local
```

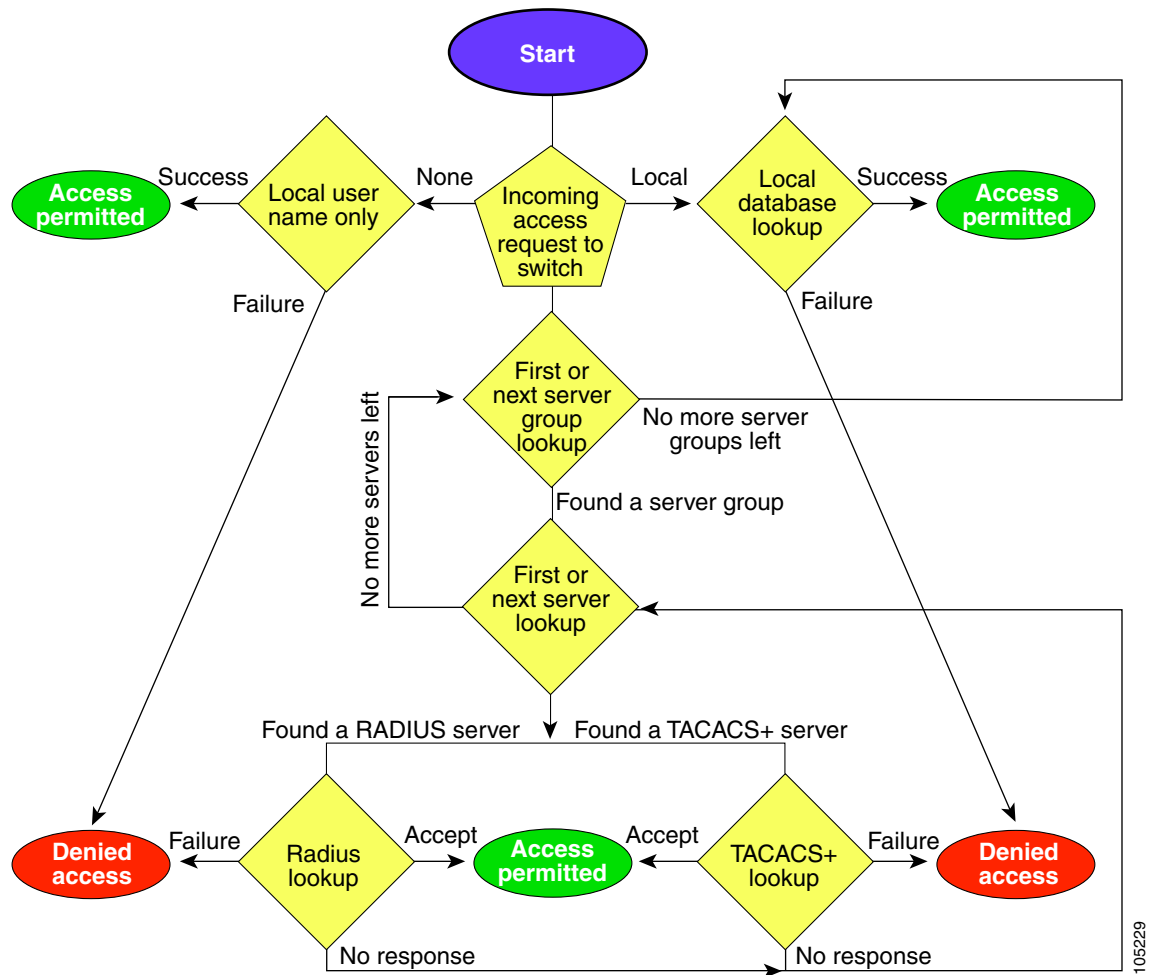
# Authentication and Authorization Process

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. The Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers or TACACS+ servers).

The following steps explain the authorization and authentication process. [Figure 16-1](#) shows a flow chart of the process.

- 
- Step 1** When you can log in to the required switch in the Cisco MDS 9000 Family, you have the option to use the Telnet, SSH, or Console login options.
- Telnet/SSH users, use the **aaa authentication login default** command.
  - Console uses, use the **aaa authentication login console** command. If this command is not configured the software automatically uses the **aaa authentication login default** command.
- Step 2** When you configure server groups using the server group authentication method, an authentication request is sent to the first AAA server in the group.
- If the AAA server fails to respond, then the next AAA server will be tried and so on until the remote server responds to the authentication request.
  - If all AAA servers in the server group fail to respond, then the servers in the next server group are tried.
  - If all configured methods fails, then local database is used for authentication.
- Step 3** When you are successfully authenticated through a remote AAA server, then the following possibilities apply:
- If AAA server protocol is RADIUS, the user roles specified in cisco-av-pair attribute is downloaded with authentication response
  - If AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for shell
  - If user roles were not retrieved successfully from remote AAA server, then the user will have role of network-operator assigned once he logs in.
- Step 4** If your user name and password are successfully authenticated, you are allowed to log in.
-

Figure 16-1 Switch Authorization and Authentication Flow



105229

# Configuring Role-Based CLI Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have permission to access that command.

To configure a new role or to modify the profile for an existing role, follow these steps:

|        | Command                                                                    | Purpose                                                                                                                                                                                                  |
|--------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                    | Enters configuration mode.                                                                                                                                                                               |
| Step 2 | switch(config)# <b>role name techdocs</b><br>switch(config-role)#          | Places you in the mode for the specified role (techdocs).<br><br><b>Note</b> The role submode prompt indicates that you are now in the role submode. This submode is now specific to the techdocs group. |
|        | switch(config)# <b>no role name techdocs</b>                               | Deletes the role called techdocs.                                                                                                                                                                        |
| Step 3 | switch(config-role)# <b>description</b><br><b>Entire Tech. Docs. group</b> | Assigns a description to the new role. The description is limited to one line and can contain spaces.                                                                                                    |
|        | switch(config-role)# <b>no description</b>                                 | Resets the description for the Tech. Docs. group.                                                                                                                                                        |

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed to perform configuration commands, and role2 users are only allowed to perform debug commands, then if Joe belongs to both role1 and role2, he can perform configuration as well as debug commands.



## Note

If you belong to multiple roles, you can execute a superset of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.



## Tip

Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

## Configuring Rules and Features for Each Role

The **rule** command specifies operations that can be performed by a specific role. Each rule consists of a rule number, a rule type (permit or deny), a command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, interface).

The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2 which is applied before rule 3 etc.

**Note**

In this case, **exec** commands refer to all commands in the EXEC mode that do not fall in the **show**, **debug**, and **clear**, categories. Up to 16 rules can be configured for each role.

To configure a new role or to modify the profile for an existing role, follow these steps:

|        | Command                                                                                                                                                                                                                                                                        | Purpose                                                                                                                                                                                                                |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                                                                                        | Enters configuration mode.                                                                                                                                                                                             |
| Step 2 | switch(config)# <b>role name sangroup</b><br>switch(config-role)#                                                                                                                                                                                                              | Places you in <i>sangroup</i> role submode.                                                                                                                                                                            |
| Step 3 | switch(config-role)# <b>rule 1 permit config</b><br>switch(config-role)# <b>rule 2 deny config</b><br><b>feature fspf</b><br>switch(config-role)# <b>rule 3 permit debug</b><br><b>feature zone</b><br>switch(config-role)# <b>rule 4 permit exec</b><br><b>feature fcping</b> | Allows users belonging to the <i>sangroup</i> role to perform all configuration commands except <b>fspf config</b> commands. They can also perform <b>zone debug</b> commands and the <b>fcping</b> EXEC mode command. |
| Step 4 | switch(config-role)# <b>no rule 4</b>                                                                                                                                                                                                                                          | Deletes rule 4 which no longer permits the <i>sangroup</i> to perform the <b>fcping</b> command.                                                                                                                       |

In Step 3, rule 1 is applied first, thus permitting all **config** commands to *sangroup* users. Rule 2 is applied next, denying FSPF configuration to *sangroup* users. As a result, *sangroup* users can perform all other **config** commands, except **fspf** configuration commands.

**Note**

The order of rule placement is important. If you had swapped these two rules and issued the **deny config feature fspf** rule first and issued the **permit config** rule next, you would be allowing all *sangroup* users to perform all configuration commands since the second rule globally overrode the first rule.

## Configuring the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE\_PKG license (see [Obtaining and Installing Licenses](#)).

You can configure a role so that it only allows commands to be performed for a selected set of VSANs. By default, the VSAN policy for any role is **permit**. In other words, the role can perform commands configured by the **rule** command in all VSANs. In order to selectively allow VSANs for a role, the VSAN policy needs to be set to **deny** and then the appropriate VSANs need to be permitted.

To configure a new role or to modify the VSAN policy for an existing role, follow these steps:

|        | Command                                                             | Purpose                                                                                                                   |
|--------|---------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                             | Enters configuration mode.                                                                                                |
| Step 2 | switch(config)# <b>role name sangroup</b><br>switch(config-role)#   | Places you in <i>sangroup</i> role submode.                                                                               |
| Step 3 | switch(config)# <b>vsan policy deny</b><br>switch(config-role-vsan) | Changes the VSAN policy of this role to <b>deny</b> and places you in a submode where VSANs can be selectively permitted. |
|        | switch(config-role)# <b>no vsan policy deny</b>                     | Deletes the configured VSAN role policy and reverts to the factory default ( <b>permit</b> ).                             |

| Step 4 | Command                                               | Purpose                                                                                                                                                        |
|--------|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
|        | switch(config-role-vsan)# <b>permit vsan 10-30</b>    | Permits this role to perform the allowed commands for VSANs 10 through 30.                                                                                     |
|        | switch(config-role-vsan)# <b>no permit vsan 15-20</b> | Removes the permission for this role to perform commands for vsan 15 to 20. So, the role is now permitted to perform commands for VSAN 10 to 14, and 21 to 30. |

**Note**

Users configured in roles where the VSAN policy set to **deny** cannot modify configuration for E ports. They can only modify configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.

**Tip**

Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN policy is set to **deny** are referred to as VSAN-restricted users. These users cannot perform the following commands that require the startup configuration to be viewed or modified:

- **copy running startup**
- **show startup**
- **show running-config diff**
- **copy startup <destination>**
- **copy <source> startup** commands.

For information on these commands, refer to [Chapter 2, “Before You Begin.”](#)

## Displaying Role-Based CLI Information

Use the **show role** command to display rules configured on the switch including those rules that have not yet been committed to persistent storage. The rules are displayed by rule number and are based on each role. All roles are displayed even if role name is not specified. See [Example 16-10](#).

### Example 16-10 Displays Information for All Roles

```
switch# show role
Role: network-admin
Description: Predefined Network Admin group. This role cannot be modified
Access to all the switch commands

Role: network-operator
Description: Predefined Network Operator group. This role cannot be modified
Access to Show commands and selected Exec commands

Role: TechDocs
vsan policy: permit (default)
```



Role: sangroup  
Description: SAN management group  
vsan policy: deny  
Permitted vsans: 10-30

| ----- |        |              |         |
|-------|--------|--------------|---------|
| Rule  | Type   | Command-type | Feature |
| ----- |        |              |         |
| 1.    | permit | config       | *       |
| 2.    | deny   | config       | fspf    |
| 3.    | permit | debug        | zone    |
| 4.    | permit | exec         | fcping  |

# Configuring CLI User Profiles

Every Cisco MDS 9000 Family switch user has related NMS information stored by the system. Your authentication information, user name, user password, password expiration date, and role membership are stored in your user profile. The CLI commands explained in this section enable you to create users and modify the profile of an existing user. These commands are restricted to privileged users as determined by your administrator.

## Creating or Updating Users

The switches use the same command (**username**) to create a user and to update an existing user. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format. By default, the user account does not expire unless you explicitly configure it to expire.



### Tip

The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nsd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys



### Note

User passwords are not displayed in the switch configuration file.

To configure a new user or to modify the profile of an existing user, follow these steps:

|        | Command                                                                          | Purpose                                                                                                                                                     |
|--------|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                          | Enters configuration mode.                                                                                                                                  |
| Step 2 | switch(config)# <b>username usam password abcd expire 2003-05-31</b>             | Creates or updates the user account (usam) along with a password (abcd) that is set to expire on 2003-05-31. The password is limited to 64 characters.      |
|        | switch(config)# <b>username msam password 0 abcd role network-operator</b>       | Creates or updates the user account (msam) along with a password (abcd) specified in clear text (indicated by 0). The password is limited to 64 characters. |
|        | switch(config)# <b>username user1 password 5 !@*asdfsdfjh!@df</b>                | Specifies an encrypted (specified by 5) password (!@*asdfsdfjh!@df) for the user account (user1).                                                           |
| Step 3 | switch(config)# <b>username usam role network-admin</b>                          | Adds the specified user (usam) to the network-admin role.                                                                                                   |
|        | switch(config)# <b>no username usam role vsan-admin</b>                          | Deletes the specified user (usam) from the vsan-admin role.                                                                                                 |
| Step 4 | switch(config)# <b>username usam sshkey fsafsd2344234234ffgsdfg</b>              | Identifies the contents of the SSH key for the specified user (usam).                                                                                       |
|        | switch(config)# <b>no username usam sshkey fsafsd2344234234ffgsdfgffsdfsfssf</b> | Deletes the SSH key content identification for the user (usam).                                                                                             |

**Note**

If the **update-snmpv3** option is used, specify the clear text and old SNMP password (see the [“Forcing Identical SNMP and CLI Passwords”](#) section on page 16-33).

## Logging out CLI Users

To log out another user on the switch, use the **clear user** command.

```
switch# clear user vsam
switch#
```

In this example, the user named vsam is logged out from the switch.

Use the **show users** command to view a list of the logged in users.

### *Example 16-11 Displays All Logged in Users*

```
switch# show users
admin pts/7 Jan 12 20:56 (10.77.202.149)
admin pts/9 Jan 12 23:29 (modena.cisco.com)
admin pts/10 Jan 13 03:05 (dhcp-171-71-58-120.cisco.com)
admin pts/11 Jan 13 01:53 (dhcp-171-71-49-49.cisco.com)
```

## Displaying User Profile Information

Use the **show user-account** command to display configured information about user accounts. See Examples [16-11](#) to [16-13](#).

### *Example 16-12 Displays Information for a Specified User*

```
switch# show user-account user1
user:user1
 this user account has no expiry date
 roles:network-operator
no password set. Local login not allowed
Remote login through RADIUS is possible
```

### *Example 16-13 Displays Information for All Users*

```
switch# show user-account
show user-account
user:admin
 this user account has no expiry date
 roles:network-admin
user:usam
 expires on Sat May 31 00:00:00 2003
 roles:network-admin network-operator
user:msam
 this user account has no expiry date
 roles:network-operator
user:user1
 this user account has no expiry date
 roles:network-operator
no password set. local login not allowed
Remote login through RADIUS is possible
```

# Configuring CLI Accounting Parameters

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

## Setting the Accounting Log Size

The **aaa accounting logsize** command sets the size limit of the accounting log file in persistent storage. The default is 15,000 bytes.

To set the log file size, follow these steps:

|        | Command                                             | Purpose                                                                       |
|--------|-----------------------------------------------------|-------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                             | Enters configuration mode.                                                    |
| Step 2 | switch(config)# <b>aaa accounting logsize 29000</b> | Sets the size of the log file on the local disk. The default is 15,000 bytes. |



### Tip

The Cisco MDS 9000 Family switch uses *Interim-Update* RADIUS accounting-request packets to communicate accounting log information to the RADIUS server. The RADIUS server must be appropriately configured to log the information communicated in these packets. Several servers typically have *Log Update/Watchdog Packets* flag in the AAA client configuration. This flag should be turned on to ensure proper RADIUS accounting.



### Note

Configuration operations are automatically recorded in the accounting log if they are performed in configuration mode. Additionally, important system events (for example, configuration save and system switchover) are also recorded in the accounting log.

## Displaying Accounting Configuration

The **show accounting** command displays configured accounting information. See Examples 16-14 to 16-16.

### Example 16-14 Displays Configured Accounting Parameters.

```
switch# show accounting config
RADIUS accounting not enabled
local accounting enabled
```

### Example 16-15 Displays Configured Log Size.

```
switch# show accounting logsize
maximum local accounting log size:29000
```

### Example 16-16 Displays the Entire Log File.

```
switch# show accounting log
```

```
Fri Jan 16 15:28:21 1981:stop:snmp_348506901_64.104.131.208:admin:
Fri Jan 16 21:17:04 1981:start:/dev/pts/0_348527824:admin:
Fri Jan 16 21:35:45 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 21:58:17 1981:start:snmp_348530297_171.71.150.105:admin:
Fri Jan 16 21:58:17 1981:stop:snmp_348530297_171.71.150.105:admin:
Fri Jan 16 21:58:18 1981:start:snmp_348530298_171.71.150.105:admin:
Fri Jan 16 21:58:18 1981:stop:snmp_348530298_171.71.150.105:admin:
...
Fri Jan 16 23:37:02 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 23:37:26 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
group:TacacsServer1
Fri Jan 16 23:45:19 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
group:TacacsServer1
Fri Jan 16 23:45:19 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
...
Fri Jan 16 23:53:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
server:Server3
Fri Jan 16 23:54:00 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
server:Server5
Fri Jan 16 23:54:22 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
server:ServerA
Fri Jan 16 23:54:25 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
server:ServerB
Fri Jan 16 23:55:03 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
...
Sat Jan 17 00:01:41 1981:start:snmp_348537701_171.71.58.100:admin:
Sat Jan 17 00:01:41 1981:stop:snmp_348537701_171.71.58.100:admin:
Sat Jan 17 00:01:42 1981:start:snmp_348537702_171.71.58.100:admin:
Sat Jan 17 00:01:42 1981:stop:snmp_348537702_171.71.58.100:admin:
...
```

# Recovering Administrator Password

An administrator can recover a password from a local console connection.

The password recovery procedure must be performed on the supervisor module that becomes the active supervisor module after the recovery procedure is completed. To ensure the other supervisor module does not become the active module, you have two options:

- Physically remove the other supervisor module from the chassis, or
- Change the other supervisor module's console prompt to the `loader>` or `switch (boot) #` prompt (see the [“Recovery from the loader> Prompt”](#) section on page 6-31) until you complete this procedure.



## Note

Password recovery is not possible from a Telnet or SSH session.

To recover a administrator's password, follow these steps:

- 
- Step 1** Reboot the switch.
- ```
switch# reload
The supervisor is going down for reboot NOW!
```
- Step 2** Press the **Ctrl-]** key sequence (when the switch begins its SAN-OS software boot sequence) to enter the `switch (boot) #` prompt (see [“Recovery Interruption”](#) section on page 6-27).
- ```
Ctrl-]
switch (boot) #
```
- Step 3** Change to configuration mode.
- ```
switch (boot) # config terminal
```
- Step 4** Enter the **admin-password** command to reset the administrator password.
- ```
switch (boot-config) # admin-password password
```
- Step 5** Exit to the EXEC mode.
- ```
switch (boot-config) # exit
switchboot#
```
- Step 6** Enter the **load** command to load the SAN-OS software.
- ```
switch (boot) # load bootflash:system.img
```
- Step 7** Save the software configuration.
- ```
switch# copy running-config startup-config
```
-

Configuring SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a host key pair. To generate a host key, use the **ssh key** command (see the “Generating an SSH Host Key Pair” section on page 16-27).

Enabling SSH Service

By default, the SSH service is disabled. To enable SSH service, issue the **ssh server enable** command. To enable or disable the SSH service, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ssh server enable updated	Enables the use of the SSH service.
	switch(config)# no ssh server enable updated	Disables (default) the use of the SSH service and resets the switch to its factory defaults.


Generating an SSH Host Key Pair

Be sure to have an SSH host key pair with the appropriate version before enabling the SSH service. The SSH service accepts three types of key pairs for use by SSH versions 1 and 2. Generate the SSH host key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

- The **rsa1** option generates the RSA1 key pair for the SSH version 1 protocol.
- The **dsa** option generates the DSA key pair for the SSH version 2 protocol.
- The **rsa** option generates the RSA key pair for the SSH version 2 protocol.

To generate the SSH host key pair, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.

	Command	Purpose
Step 2	switch(config)# ssh key rsa1 1024 generating rsa1 key..... generated rsa1 key switch(config)#	Generates the RSA1 host key pair.
	switch(config)# ssh key dsa 1024 generating dsa key..... generated dsa key switch(config)#	Generates the DSA host key pair.
	switch(config)# ssh key rsa 1024 generating rsa key..... generated rsa key switch(config)#	Generates the RSA host key pair.
	switch(config)# no ssh key rsa 1024 cleared RSA keys switch(config)#	Clears the RSA host key pair configuration.
		Caution If you delete all of the SSH keys, you cannot start a new session.

Using the force Option

If the SSH key pair option is already generated for the required version, use the **force** option to overwrite the previously generated key pair.

To overwrite the previously generated key pair, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ssh key dsa 768 ssh key dsa 512 dsa keys already present, use force option to overwrite them	Tries to set the host key pair. If a required host key pair is already configured, use the force option to overwrite that host key pair.
	switch(config)# ssh key dsa 512 force deleting old dsa key..... generating dsa key..... generated dsa key switch(config)#	

Clearing SSH Hosts

To manually clear trusted SSH host entries, issue the **clear ssh hosts** command at the switch prompt:

Example 16-17 Clearing Configured SSH Hosts

```
switch# clear ssh hosts
switch#
```

This command clears reset-reason information from NVRAM and volatile storage.

Displaying SSH Protocol Status

Use the **show ssh server** command to display the status of the SSH protocol (enabled or disabled) and the versions that are enabled for that switch. See [Example 16-18](#).

Example 16-18 Displays SSH Protocol Status

```
switch# show ssh server
ssh is enabled
version 1 enabled
version 2 enabled
```

Use the **show ssh key** command to display the host key pair details for the specified key or for all keys, if no key is specified. See [Example 16-19](#).

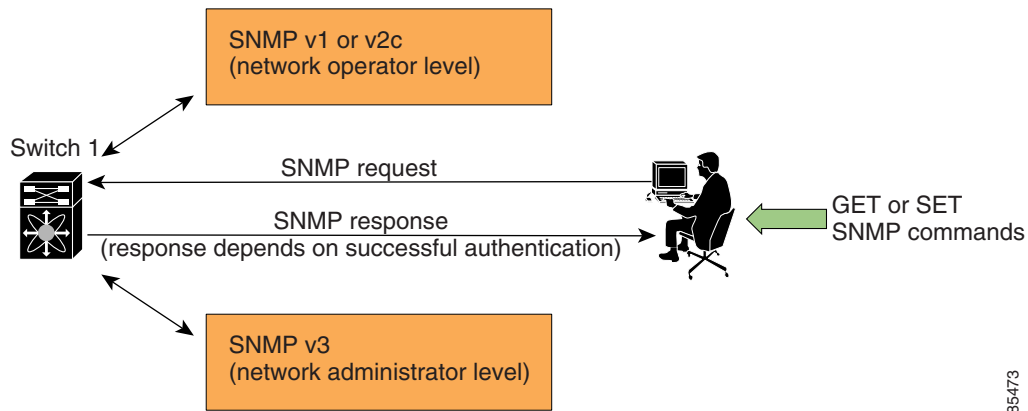
Example 16-19 Displays Host Key Pair Details

```
switch# show ssh key
rsa1 Keys generated:Sun Jan 13 07:16:26 1980
1024 35
fingerprint:
1024 67:76:02:bd:3e:8d:f5:ad:59:5a:1e:c4:5e:44:03:07
could not retrieve rsa key information
dsa Keys generated:Sun Jan 13 07:40:08 1980
ssh-dss AAAAB3NzaC1kc3MAAABBAJTCRQOydNRel2v7uiO6Fix+OTn8eGdnnDVxw5eJs5OcOEXOyjaW
cMMYsEgxc9ada1NElp8Wy7GPMWGOQYj9CU0AAAAMCcWhNN18zFNOIPo7cU3t7d0iEbAAAAQBdQ8UAO
i/Cti84qFb3kTqXlS9mEhdQUo0lHcH5bw5PKfj2Y/dLR437zCBKXetPj4p7mhQ6Fq5os8RZtJEyOsNsA
AABAA0oxZbPyWeR5NHATXiyXdPI7j9i8fgyn9FNipMkOF2Mn75Mi/lqQ4NIq0gQNvQOx27uCeQlRts/Q
wI4q68/eaw==
fingerprint:
512 f7:cc:90:3d:f5:8a:a9:ca:48:76:9f:f8:6e:71:d4:ae
```

SNMP Security

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3 (see Figure 16-2).

Figure 16-2 SNMP Security



85473



Note

Users and roles configured through the CLI are different from users and roles configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

SNMP users are different from CLI users. SNMP users also have role-based authentication for roles and authorization purposes.

SNMP Version 1 and Version 2c

SNMPv1 and SNMPv2c use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

SNMP Version 3

SNMPv3 is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP Access Control Lists (IP-ACLs). See the [“IP Access Control Lists” section on page 20-5](#).

Group-Based SNMP Access

**Note**

Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once the your user name is created, your roles are set up by your administrator, and you are added to the roles.

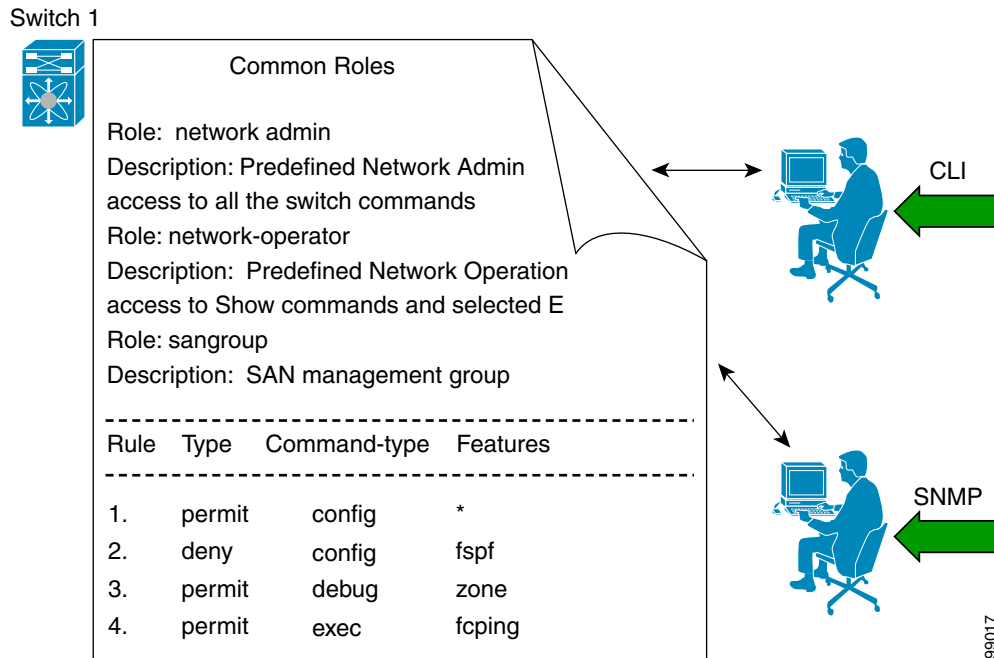
**Note**

Users configured through the CLI are different from users configured through SNMP. These configurations do not directly correspond with each other. However, you can configure both CLI and SNMP identically, if required.

Configuring Common Roles

From Release 1.2(x), CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles. You can use SNMP to modify a role that was created using CLI and vice versa (see [Figure 16-3](#)).

Figure 16-3 Common Roles



Each role in SNMP is the same as a role created or modified through the CLI (see “[Configuring Role-Based CLI Authorization](#)” section on page 16-18).

Each role can be restricted to one or more VSAN as required.

You can create new roles or modify existing roles using SNMP or the CLI.

- **SNMP**—Use the CISCO-COMMON-ROLES-MIB to configure or modify roles. Refer to the *Cisco MDS 9000 Family MIB Reference Guide* for more information.
- **CLI**—Use the **role name** command.

Creating and Modifying Users

You can create users or modify existing users using SNMP or the CLI.

- **SNMP**—Create a user as a clone of an existing user in the vsmUserTable on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC2574.



Note

You must explicitly configure password(s) for SNMP users. The SNMP user passwords are not generated as the part of the configuration file as they are not portable across devices. The password is limited to a minimum of 8 characters and a maximum of 64 characters.

**Tip**

An SNMP user must be created on each switch to which the user requires access. If the user is managing 10 switches, each of the 10 switches must have the SNMP user defined.

- CLI—You can create a user or modify an existing user using the **snmp-server user** command.

By default only two roles are available in a Cisco MDS 9000 Family switch—network-operator and network-admin. You can also use any role that is configured in the Common Roles database (see the “Configuring Common Roles” section on page 16-32).

To create or modify SNMP users using the CLI, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# snmp-server user joe network-admin auth sha abcd1234	Creates or modifies the settings for a user (joe) in the network-admin role using the HMAC-SHA-96 authentication password (abcd1234).
	switch(config)# snmp-server user sam network-admin auth md5 abcdefgh switch(config)#	Creates or modifies the settings for a user (sam) in the network-admin role using the HMAC-MD5-96 authentication password (abcdefgh).
	switch(config)# snmp-server user Bill network-admin auth sha abcd1234 priv abcdefgh	Creates or modifies the settings for a user (network-admin) in the network-admin role using the HMAC-SHA-96 authentication level and privacy encryption parameters.
	switch(config)# no snmp-server user usernameA	Deletes the user (usernameA) and all associated parameters.
	switch(config)# snmp-server user user1 network-admin auth md5 0xab0211gh priv 0x45abf342 localizedkey	Specifies the password to be in localized key format (see RFC2574). The localized key is provided in the hex format (for example, 0xacbdef).

**Note**

Avoid using the **localizedkey** option when configuring an SNMP user from CLI. The localized keys are not portable across devices as they contain device engine ID information. If a configuration file is copied to the device, the passwords may not be set correctly if the configuration file was generated at a different device. Explicitly configure the desired passwords after copying the configuration into the device.

Forcing Identical SNMP and CLI Passwords

You can force the SNMPv3 password and the CLI password to be the same. You must know the SNMPv3 password to change the password using the CLI. Use the CLI password to synchronize the SNMP password. The password is limited to a minimum of 8 characters and a maximum of 64 characters.

**Caution**

To change the SNMP password, a clear text CLI password is required.

To modify the secret key for an SNMPv3 user, refer to RFC2574.

To update the SNMPv3 password from the CLI, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# username joe password wxyz6789 update-snmpv3 abcd1234	Updates the SNMPv3 password for the specified user (joe). The local CLI password and the SNMP password are updated. If user Joe does not exist, the command fails.

Assigning Users to Roles

Once the user and the role are created, the administrator should configure an entry in the vacmSecurityToGroupTable to add the configured user to a configured role.

To assign users to roles through SNMP, refer to RFC2575.

To assign users to roles through the CLI, refer to the procedure specified in the [“Creating and Modifying Users”](#) section on page 16-32.

Adding or Deleting Communities

You can configure read-only or read-write access for SNMP users by using the **snmp-server community** CLI command. Use the **no** form of the command to delete the configured community. Refer to RFC2576.

To create an SNMPv1 or SNMPv2c community, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# snmp-server community snmp_Community ro	Adds read-only access for the specified SNMP community.
	switch(config)# snmp-server community snmp_Community rw	Adds read-write access for the specified SNMP community.
	switch(config)# no snmp-server community snmp_Community	Deletes access for the specified SNMP community (default).

Displaying SNMP Security Information

Use the **show snmp** commands to display configured SNMP information (see [Example 16-20](#) and [16-22](#)).

Example 16-20 Displays SNMP User Details

```
switch# show snmp user
```

User	Group	Auth	Priv
steve	network-admin	md5	des
sadmin	network-admin	md5	des
stever	network-operator	md5	des

Example 16-21 Displays SNMP Community Information

```
switch# show snmp community
```

Community	Access
private	rw
public	ro
v93RACqPNH	ro

Example 16-22 Displays SNMP Host Information

```
switch# show snmp host
```

Host	Port	Version	Level	Type	SecName
171.16.126.34	2162	v2c	noauth	trap	public
171.16.75.106	2162	v2c	noauth	trap	public
...					
171.31.58.97	2162	v2c	auth	trap	public
...					

Displaying SNMP Counter Information

The **show snmp** command displays counter information for SNMP contact, location, and packet settings. This command provides information that is used entirely by the Cisco MDS 9000 Family Fabric Manager (refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*). See [Example 16-23](#).

Example 16-23 Displays SNMP

```
switch# show snmp
switch# show snmp
sys contact:
sys location:
1631 SNMP packets input
    0 Bad SNMP versions
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
64294 Number of requested variables
1 Number of altered variables
1628 Get-request PDUs
0 Get-next PDUs
1 Set-request PDUs
152725 SNMP packets output
```

```

0 Too big errors
1 No such name errors
0 Bad values errors
0 General errors
Community
-----
public
User
-----
admin

Access
-----
rw
Group
-----
network-admin

Auth Priv
-----
md5 no

```

Default Security Settings

Table 16-1 lists the default settings for all security features in any switch.

Table 16-1 Default Security Settings

Parameters	Default
Roles in each switch (for CLI and SNMP users)	Two default roles—network-operator and network-admin.
AAA authentication login	Local authentication is enabled. If the Telnet or SSH options are not specified, the command applies to both.
Telnet server	Enabled.
Accounting log file size on local disk	15,000 bytes.
User's account expiration	Does not expire unless you explicitly configure it to expire.
User name	admin.
User password	admin.
Configured RADIUS sever	Allows access to all RADIUS servers.
RADIUS server timeout interval	The default time-out is one (1) seconds.
RADIUS preshared key	No key is configured.
RADIUS key encryption	clear text (0)—not encrypted.
RADIUS server connection attempts	A switch tries to connect to a RADIUS server once (1).
RADIUS Authentication port	UDP port 1812.
RADIUS Accounting port	UDP port 1813.
Server key encryption	clear text (0)—not encrypted.
TACACS+	Disabled
Configured TACACS+ sever	Allows access to all TACACS+ servers.
TACACS+ server timeout interval	The default time-out is one (5) seconds.
TACACS+ preshared key	No key is configured.
TACACS+ key encryption	clear text (0)—not encrypted
TACACS+ server connection attempts	A switch tries to connect to a TACACS+ server once (1).
TACACS+ Authentication port	UDP port 49.
VSAN policy	Permit.



Configuring Fabric Security

Fibre Channel Security Protocol (FC-SP) capabilities in SAN-OS 1.3(x) provides switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol implemented in SAN-OS 1.3(x) to provide authentication between Cisco MDS switches and other devices. It consists of the CHAP protocol combined with the Diffie-Hellman exchange.

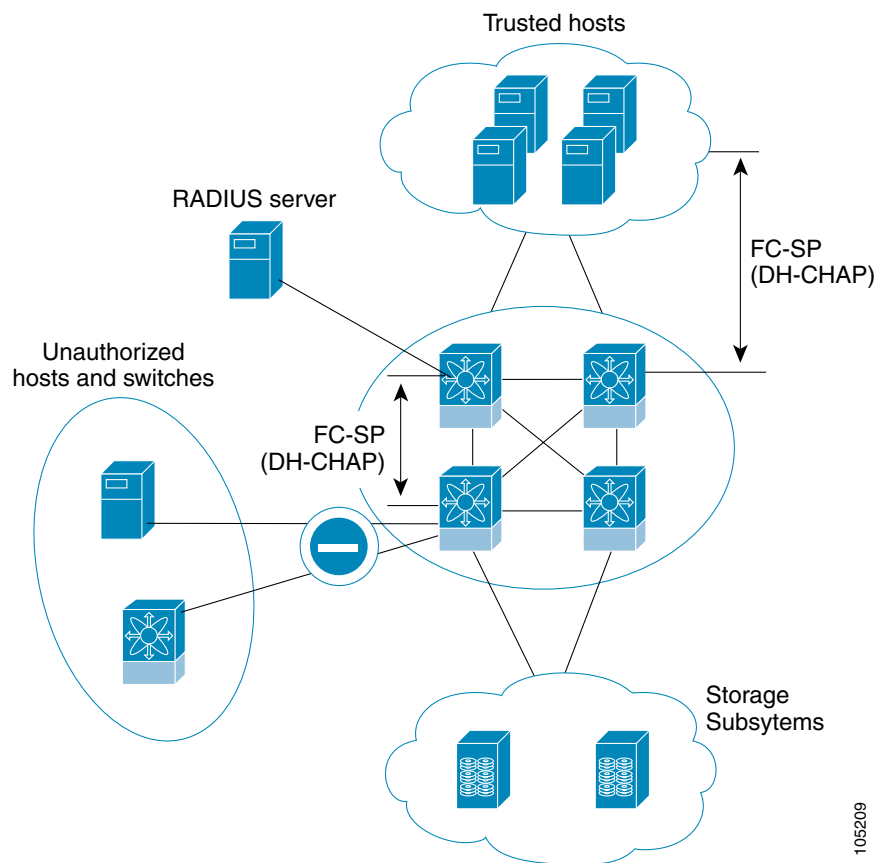
This chapter includes the following sections:

- [About Fabric Authentication, page 17-2](#)
- [About DHCHAP, page 17-3](#)
- [DHCHAP Compatibility with Existing MDS Features, page 17-3](#)
- [Configuring DHCHAP Authentication, page 17-3](#)
- [Enabling DHCHAP, page 17-4](#)
- [Configuring DHCHAP Authentication Modes, page 17-4](#)
- [Configuring the DHCHAP Hash Algorithm, page 17-5](#)
- [Configuring DHCHAP Passwords, page 17-6](#)
- [Configuring Passwords for Other Devices, page 17-8](#)
- [Configuring the DHCHAP Timeout Value, page 17-8](#)
- [Displaying Protocol Security Information, page 17-9](#)
- [DHCHAP AAA Authentication, page 17-10](#)
- [Default Fabric Security Settings, page 17-10](#)

About Fabric Authentication

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switches and hosts authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands, cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in inter-switch link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family (see [Figure 17-1](#)).

Figure 17-1 Authentication between Switches and Hosts



105209



Note

Fibre Channel (FC) host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

About DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, thus preventing unauthorized devices from accessing the switch.

**Note**

The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD-5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license (see [Obtaining and Installing Licenses](#)).

DHCHAP Compatibility with Existing MDS Features

This section identifies the impact of configuring the DHCHAP feature along with existing MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with FCIP interface just as it would with a physical interface.
- Port security or fabric binding—Fabric binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on per-VSAN basis.
- High availability—DHCHAP authentication works transparently with existing HA features.

Configuring DHCHAP Authentication

To configure DHCHAP authentication using the local password database, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Enable DHCHAP (see the “Enabling DHCHAP” section on page 17-4). |
| Step 2 | Identify and configure the DHCHAP authentication modes (see the “Configuring DHCHAP Authentication Modes” section on page 17-4). |
| Step 3 | Configure the hash algorithm and DH group (see the “Configuring the DHCHAP Hash Algorithm” section on page 17-5). |
| Step 4 | Configure the password for the local switch and other switches in the fabric (see the “Configuring DHCHAP Passwords” section on page 17-6). |
| Step 5 | Configure the timeout value for reauthentication (see the “Configuring the DHCHAP Timeout Value” section on page 17-8). |
| Step 6 | Verify the DHCHAP configuration (see the Displaying Protocol Security Information , page 17-9). |
-

Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.


To enable DHCHAP for a Cisco MDS switch, follow these steps:

	Command	Purpose
Step 1	switch# <code>config t</code>	Enters configuration mode.
Step 2	switch(config)# <code>fcsp enable</code>	Enables the DHCHAP in this switch.
	switch(config)# <code>no fcsp enable</code>	Disables (default) the DHCHAP in this switch.

Configuring DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode. When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to an isolated state.
- Auto-Active—During switch initialization if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—Does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.



Note

Whenever DHCHAP port mode is changed to a mode other than the **Off** mode, reauthentication is performed.

Table 17-2 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Table 17-1 DHCHAP Authentication Status Between Two MDS Switches

Switch N	Switch 1			
DHCHAP Modes	on	auto-active	auto-passive	off
on	FC-SP authentication is performed	FC-SP authentication is performed	FC-SP authentication is performed	Link is brought down
auto-Active			FC-SP authentication is <i>not</i> performed.	FC-SP authentication is <i>not</i> performed.
auto-Passive				
off	Link is brought down	FC-SP authentication is <i>not</i> performed.		

To enable the DHCHAP mode for a particular interface, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc2/1-3 switch(config-if)#	Enters the interface submenu.
Step 3	switch(config-if)# fcsp on	Sets the DHCHAP mode for the selected interfaces to be in the on state.
	switch(config-if)# no fcsp on	Reverts to the factory default of auto-passive for these three interfaces.
Step 4	switch(config-if)# fcsp auto-active 0	Changes the DHCHAP authentication mode for the selected interfaces to auto-active . The 0 indicates that the port does not perform authentication
	switch(config-if)# fcsp auto-active 120	Changes the DHCHAP authentication mode to auto-active for the selected ports to reauthenticate every two hours (120 minutes) after the initialization authentication.
	switch(config-if)# fcsp auto-active	Changes the DHCHAP authentication mode to auto-active for the selected ports.

Configuring the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD-5 followed by SHA-1 for DHCHAP authentication.



Tip

If you change the hash algorithm configuration, ensure to change it globally for all switches in the fabric.



Caution

RADIUS and TACACS+ protocols always use MD-5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

To change the hash algorithm, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# fcsp dhchap hash sha1	Configures the use of only the SHA-1 hash algorithm.
	switch(config)# fcsp dhchap hash MD5	Configures the use of only the MD-5 hash algorithm.
	switch(config)# fcsp dhchap hash md5 sha1	Defines the use of the default hash algorithm priority list of MD-5 followed by SHA-1 for DHCHAP authentication.
	switch(config)# no fcsp dhchap hash sha1	Reverts to the factory default priority list of the MD-5 hash algorithm followed by the SHA-1 hash algorithm.

Configuring DHCHAP Groups

All switches in the Cisco MDS Family support all DHCHAP groups specified in the standard: 0 (null DH group which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.



Tip

If you change the DH group configuration, ensure to change it globally for all switches in the fabric.

To change the DHCHAP settings, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# fcsp dhchap group 2 3 4	Prioritizes the use of DH group 2, 3, and 4 in the configured order.
	switch(config)# no fcsp dhchap group 0	Reverts to the DHCHAP factory default order of 0, 4, 1, 2, and 3 respectively.

Configuring DHCHAP Passwords

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three approaches to manage passwords for all switches in the fabric which participate in DHCHAP:

- Approach 1: Use the same password for all switches in the fabric—the simplest approach. When you add a new switch, you will use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from outside maliciously attempts to access any one switch in the fabric
- Approach 2: Use a different password for each switch and maintain that password list in each switch in the fabric—when you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.

- Approach 3: Use different passwords for different switches in the fabric—when you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This approach requires considerable password maintenance by the user.

**Tip**

We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use local password database, you can continue to do so using Approach 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database. Refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for further information.

**Note**

All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

To configure the DHCHAP password for the local switch, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# fcsp dhchap password 0 mypassword	Configures a clear text password for the local switch.
	switch(config)# fcsp dhchap password 0 mypassword 30:11:bb:cc:dd:33:11:22	Configures a clear text password for the local switch to be used for the device with the specified WWN.
	switch(config)# no fcsp dhchap password 0 mypassword 30:11:bb:cc:dd:33:11:22	Removes the clear text password for the local switch to be used for the device with the specified WWN.
	switch(config)# fcsp dhchap password 7 sfsfdf	Configures a password entered in an encrypted format for the local switch.
	switch(config)# fcsp dhchap password 7 sfsfdf 29:11:bb:cc:dd:33:11:22	Configures a password entered in an encrypted format for the local switch to be used for the device with the specified WWN.
	switch(config)# no fcsp dhchap password 7 sfsfdf 29:11:bb:cc:dd:33:11:22	Removes the password entered in an encrypted format for the local switch to be used for the device with the specified WWN.
	switch(config)# fcsp dhchap password mypassword1	Configures a clear text password for the local switch to be used with any connecting device.

Configuring Passwords for Other Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).



Note

The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

To locally configure the device name of another switch in the fabric, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code>	Enters configuration mode.
Step 2	<code>switch(config)# fcsp dhchap devicename 00:11:22:33:44:aa:bb:cc password NewPassword</code>	Configures a password for another switch in the fabric which is identified by the Switch WWN device name.
	<code>switch(config)# no fcsp dhchap devicename 00:11:22:33:44:aa:bb:cc password NewPassword</code>	Removes the password entry for this switch from the local authentication database.
	<code>switch(config)# fcsp dhchap devicename 00:11:55:66:00:aa:bb:cc password 0 NewPassword</code>	Configures a clear text password for another switch in the fabric which is identified by the Switch WWN device name.
	<code>switch(config)# fcsp dhchap devicename 00:11:22:33:55:aa:bb:cc password 7 asdfklkj</code>	Configures a password entered in an encrypted format for another switch in the fabric which is identified by the Switch WWN device name.

Configuring the DHCHAP Timeout Value

During the DHCHAP protocol exchange if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured all switches in the fabric.

To change the hash algorithm, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code>	Enters configuration mode.
Step 2	<code>switch(config)# fcsp timeout 60</code>	Configures the reauthentication timeout to be 60 seconds.
	<code>switch(config)# no fcsp timeout 60</code>	Reverts to the factory default of 30 seconds.

Displaying Protocol Security Information

Use the **show fcsp** commands to display configurations for the local database (see [Example 17-1](#) and [17-6](#)).

Example 17-1 *Displays DHCHAP Configurations in FC Interfaces*

```
switch# show fcsp interface fc1/9

fc1/9:
  fcsp authentication mode:SEC_MODE_ON
  Status: Successfully authenticated
```

Example 17-2 *Displays DHCHAP Statistics for a FC Interfaces*

```
switch# show fcsp interface fc1/9 statistics

fc1/9:
  fcsp authentication mode:SEC_MODE_ON
  Status: Successfully authenticated
  Statistics:
  FC-SP Authentication Succeeded:5
  FC-SP Authentication Failed:0
  FC-SP Authentication Bypassed:0
```

Example 17-3 *Displays the FC-SP WWN of the Device Connected through a Specified Interface*

```
switch# show fcsp interface fc 2/1 wwn

fc2/1:
  fcsp authentication mode:SEC_MODE_ON
  Status: Successfully authenticated
  Other device's WWN:20:00:00:e0:8b:0a:5d:e7
```

Example 17-4 *Displays Hash Algorithm and DHCHAP Groups Configured for the Local Switch*

```
switch# show fcsp dhchap
Supported Hash algorithms (in order of preference):
DHCHAP_HASH_MD5
DHCHAP_HASH_SHA_1

Supported Diffie Hellman group ids (in order of preference):
DHCHAP_GROUP_NULL
DHCHAP_GROUP_1536
DHCHAP_GROUP_1024
DHCHAP_GROUP_1280
DHCHAP_GROUP_2048
```

Example 17-5 *Displays the DHCHAP Local Password Database*

```
switch# show fcsp dhchap database
DHCHAP Local Password:
  Non-device specific password:mypassword1
  Password for device with WWN:29:11:bb:cc:dd:33:11:22 is pjoalf
  Password for device with WWN:30:11:bb:cc:dd:33:11:22 is mypassword

Other Devices' Passwords:
  Password for device with WWN:00:11:22:33:44:aa:bb:cc is NewPassword
```

Example 17-6 Displays the ASCII Representation of the Device WWN

```
switch# show fcsp asciiwwn 30:11:bb:cc:dd:33:11:22
Ascii representation of WWN to be used with AAA servers:0x_3011bbccdd331122
```



Tip

Use the ASCII representation of the Device WWN (identified in bold in [Example 17-6](#)) to configure the switch information on RADIUS and TACACS+ servers.

DHCHAP AAA Authentication

You can individually set authentication options using the **aaa authentication dhchap** command. If authentication is not configured, local authentication is used by default.

To configure the AAA authentication, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# aaa authentication dhchap default group TacacsServer1	Enables DHCHAP to use the TACACS+ server group (in this example, TacacsServer1) for authentication.
	switch(config)# aaa authentication dhchap default local	Enables DHCHAP for local authentication.
	switch(config)# aaa authentication dhchap default group RadiusServer1	Enables DHCHAP to use the RADIUS server group (in this example, RadiusServer1) for authentication.

Default Fabric Security Settings

[Table 17-2](#) lists the default settings for all fabric security features in any switch.

Table 17-2 Default Fabric Security Settings

Parameters	Default
DHCHAP feature	Disabled.
DHCHAP hash algorithm	A priority list of MD-5 followed by SHA-1 for DHCHAP authentication
DHCHAP authentication mode	auto-passive.
DHCHAP group default priority exchange order	0, 4, 1, 2, and 3 respectively.
DHCHAP timeout value	30 seconds.



Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and reports these intrusions to the administrator.

This chapter includes the following sections:

- [Port Security Features, page 18-2](#)
- [About Auto-Learn, page 18-2](#)
- [Manually Configuring Port Security, page 18-5](#)
- [Copying the Port Security Database, page 18-8](#)
- [Database Scenarios, page 18-8](#)
- [Deleting the Port Security Database, page 18-10](#)
- [Displaying Port Security Commands, page 18-10](#)
- [Default Port Security Settings, page 18-12](#)



Note

Port security is only supported for Fibre Channel ports.

Port Security Features

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through syslog messages.

Enforcing Port Security

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up (**no shutdown** command).

The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learn

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. The **auto-learn** option allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature to activate port security feature for the first time as it saves tedious manual configuration for each port. Auto-learn is configured on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learnt, even if you have not configured any port access. Learnt entries on a port are cleaned up after a **shutdown** command is issued on that port.

Activating Port Security

By default, the port security feature is not activated.

To enable the port security feature, follow these steps:

	Command	Purpose
Step 1	switch# confi t switch(config)#	Enters configuration mode.

	Command	Purpose
Step 2	<code>switch(config)# port-security activate vsan 1</code>	Activates the port security database for the specified VSAN, and automatically enables auto-learn.
	<code>switch(config)# no port-security activate vsan 1</code>	Deactivates the port security database for the specified VSAN, and automatically disables auto-learn.

When you activate the port security feature, the **auto-learn** option is also automatically enabled. You can choose to activate the port-security feature and disable **auto-learn** using the **port-security activate vsan number no-auto-learn** command. In this case, you need to manually populate the port security database by individually securing each port.

To enable the port security feature, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# port-security activate vsan 1 no-auto-learn</code>	Disables the auto-learn feature for the port security database in VSAN 1.

Configuring Auto-Learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, the **auto-learn** option is disabled by default.
- If the port security feature is activated, the **auto-learn** option is enabled by default (unless it is turned off using the **port-security activate vsan number no-auto-learn** command).

To enable the auto-learn option, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# port-security auto-learn vsan 1</code>	Enables auto-learning so the switch can learn about any device that is allowed to access VSAN 1. These devices are logged in the port security active database.
	<code>switch(config)# no port-security auto-learn vsan 1</code>	Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learnt up to this point.



Tip

If the **auto-learn** option is enabled on a VSAN, you cannot activate the database for that VSAN without the **force** option.

Table 18-1 summarizes the authorized connection for device requests.

Table 18-1 Auto-learn Device Authorization

Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization	Condition
Configured with one or more switch ports	a switch on configured ports	Permitted	1
	a switch on other ports	Denied	2
Not configured	a port that is not configured	Permitted if auto-learn enabled	3
		Denied if auto-learn disabled	4
Configured or not configured	a switch port that allows any device	Permitted	5
Configured to login to any switch port	any port on the switch	Permitted	6
Not configured	a port configured with some other device	Denied	7

Authorization Scenario

Assuming that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc1/1 (F1)
- A pWWN (P2) is allowed access through interface fc1/1 (F1)
- A nWWN (N1) is allowed access through interface fc1/2 (F2)
- Any WWN is allowed access through interface fc1/3 (F3)
- A nWWN (N3) is allowed access through any interface
- A pWWN (P3) is allowed access through interface fc1/4 (F4)
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13)
- A pWWN (P10) is allowed access through interface fc1/11 (F11)

Table 18-2 summarizes the port security authorization results for this active database.

Table 18-2 Authorization Results for Scenario

Scenario	Device Connection Request	Authorization	Condition	Reason
1	P1, N2, F1	Permitted	1	No conflict
2	P2, N2, F1	Permitted	1	No conflict
3	P3, N2, F1	Denied	2	F1 is bound to P1/P2
4	P1, N3, F1	Permitted	6	Wildcard match for N3
5	P1, N1, F3	Permitted	5	Wildcard match for F3
6	P1, N4, F5	Denied	2	P1 is bound to F1
7	P5, N1, F5	Denied	2	N1 is only allowed on F2

Table 18-2 Authorization Results for Scenario (continued)

Scenario	Device Connection Request	Authorization	Condition	Reason
8	P3, N3, F4	Permitted	1	No conflict
9	S1, F10	Permitted	1	No conflict
10	S2, F11	Denied	7	P10 is bound to F11
11	P4, N4, F5 (auto-learn on)	Permitted	3	No conflict
12	P4, N4, F5(auto-learn off)	Denied	4	No match
13	S3, F5 (auto-learn on)	Permitted	3	No conflict
14	S3, F5 (auto-learn off)	Denied	4	No match
15	P1, N1, F6 (auto-learn on)	Denied	2	P1 is bound to F1
16	P5, N5, F1 (auto-learn on)	Denied	7	P3 is bound to F1
17	S3, F4 (auto-learn on)	Denied	7	P3 paired with F4
18	S1, F3 (auto-learn on)	Permitted	5	No conflict
19	P5, N3, F3	Permitted	6	Wildcard match for F3 and N3
20	P7, N3, F9	Permitted	6	Wildcard match for N3

Manually Configuring Port Security

To configure port security in any switch in the Cisco MDS 9000 Family, follow these steps:

-
- Step 1** Identify the WWN of the ports that need to be secured (see the [“Identifying WWNs to Configure Port Security” section on page 18-5](#)).
- Step 2** Secure the fWWN to an authorized nWWN or pWWN (see the [Securing Authorized Ports, page 18-6](#)).
- Step 3** Activate the port security database (see the [Activating the Port Security Database, page 18-6](#)).
- Step 4** Verify your configuration (see the [Displaying Port Security Commands, page 18-10](#))
-

Identifying WWNs to Configure Port Security

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or the fWWN.
- Identify devices by the pWWN or nWWN.
- If an Nx port:
 - is allowed to login to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
 - nWWN is bound to a Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.
- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.

- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.
- Saving the running configuration (using the **copy running start** command) saves the configuration database and activated entries in the active database. Learnt entries in the active database are not saved.

Securing Authorized Ports

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

To configure port security, follow these steps:

	Command	Purpose
Step 1	switch# confi g t switch(config)#	Enters configuration mode.
Step 2	switch(config)# port-security database vsan 1 switch(config-port-security)#	Enters the port security database mode for the specified VSAN.
	switch(config)# no port-security database vsan 1 switch(config)#	Deletes the port security configuration database from the specified VSAN.
Step 3	switch(config-port-security)# swwn 20:01:33:11:00:2a:4a:66 interface port-channel 5	Configures the specified sWWN to only login through PortChannel 5.
	switch(config-port-security)# any-wwn interface fc1/1 - fc1/8	Configures any WWN to login through the specified interfaces.
	switch(config-port-security)# pwwn 20:11:00:33:11:00:2a:4a fwwn 20:81:00:44:22:00:4a:9e	Configures the specified pWWN to only log in through the specified fWWN.
	switch(config-port-security)# no pwwn 20:11:00:33:11:00:2a:4a fwwn 20:81:00:44:22:00:4a:9e	Deletes the specified pWWN configured in the previous step.
	switch(config-port-security)# nwwn 26:33:22:00:55:05:3d:4c fwwn 20:81:00:44:22:00:4a:9e	Configures the specified nWWN to log in through the specified fWWN.
	switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66	Configures the specified pWWN to login through any port on the local switch.
	switch(config-port-security)# any-wwn interface fc3/1	Configures any WWN to log in through the specified interface.
	switch(config-port-security)# no any-wwn interface fc2/1	Deletes the wildcard configured in the previous step.

Activating the Port Security Database

When you activate the port security database, all entries in the configured database are copied to the active database. After the database is activated, subsequent device login is subject to the activated port bound WWN pairs. Additionally, all devices that have already logged into the VSAN at the time of activation are also learnt and added to the active database. If the **auto-learn** option is already enabled in a VSAN, you will not be allowed to activate the database (see the [About Auto-Learn, page 18-2](#)).

To activate the port database, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# port-security activate vsan 1 switch(config-port-security)#	Activates the port security database for the specified VSAN, and automatically turns on the auto-learn feature.
	switch(config)# no port-security activate vsan 1	Deactivates port security, deletes the active database, and disables auto-learn.

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database. View such entries using the **port-security database diff active vsan** command.
- The **auto-learn** option was enabled before the activation. See the [“Forcing Port Security Activation” section on page 18-7](#) to reactivate a database in this state.
- The exact security is not configured for each PortChannel member.
- If the configured database is empty and the active database is not.

Forcing Port Security Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the **force** option.



Note

An activation using the **force** option logs out existing devices if they violate the active database.

To forcefully activate the port security database, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# port-security activate vsan 1 force	Forces the VSAN 1 port security database to activate despite conflicts.
	switch(config)# no port-security activate vsan 1 force	Reverts to the previously-configured state or to the factory default (if no state is configured).

Reactivating the Database

If the **auto-learn** option is enabled and you activate the database, you will not be allowed to proceed. To reactivate the database, follow these steps:

- Step 1

Disable the **auto-learn** option (the **no port-security auto-learn vsan number** command).
- Step 2

Copy the active database to the configured database (the **port-security database copy vsan** command). This command will overwrite the configuration database with the active database.
- Step 3

Activate the database using the **port-security activate vsan number** command.

Copying the Port Security Database

Use the **port-security database copy vsan** command to copy from the active to the configured database. If the active database is empty, this command is not accepted.

```
switch# port-security database copy vsan 1
switch#
```

Use the **port-security database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.

```
switch# port-security database diff active vsan 1
```

Use the **port-security database diff config vsan** command to obtain information on the differences between the configuration database and the active database.

```
switch# port-security database diff config vsan 1
```

Database Scenarios

Table 18-3 lists the differences and interaction between the active and configuration databases.

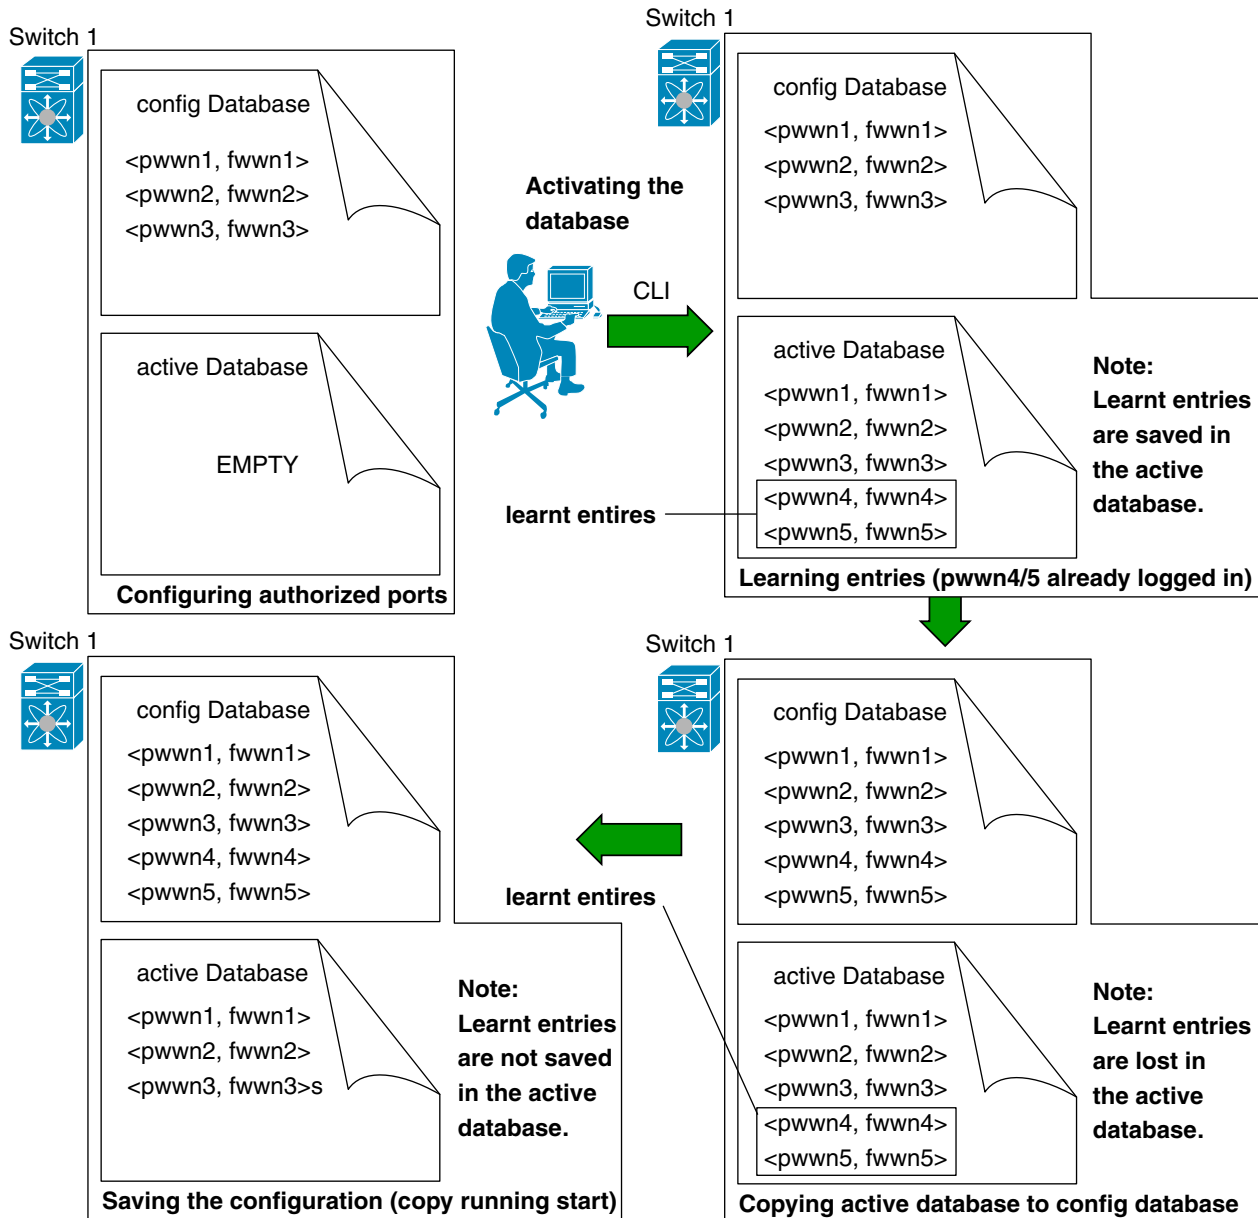
Table 18-3 Active and Configuration Port Security Databases

Configuration Database	Active Database
Read-write.	Read only.
Saving the configuration saves all the entries in the configuration database.	Saving the configuration only saves the activated entries. Learnt entries are not saved.
Once activated, the configuration database can be modified without any effect on the active database.	Once activated, all devices that have already logged into the VSAN are also learnt and added to the active database.
You can overwrite the configuration database with the active database using the port-security database copy vsan command).	You can overwrite the active database with the configured database by activating the port security database. An activation using the force option may violate the entries already configured in the active database.

The **port-security database diff active vsan** command lists the differences between the active database and the configuration database.

Figure 18-1 depicts various scenarios to depict the active database and the configuration database status based on port security configurations.

Figure 18-1 Port Security Database Scenarios



99301

Clearing the Port Security Database

Use the **clear port-security statistics** command to clear all existing statistics from the port security database for a specified VSAN.

```
switch# clear port-security statistics vsan 1
```

Use the **clear port-security database auto-learn interface** command to clear any learnt entries in the active database for a specified interface within a VSAN. The active database is read-only and this command can be used when resolving conflicts.

```
switch# clear port-security database auto-learn interface fc1/1 vsan 1
```

Use the **clear port-security database auto-learn** command to clear any learnt entries in the active database up to for the entire VSAN. The active database is read-only and this command can be used when resolving conflicts.

```
switch# clear port-security database auto-learn vsan 1
```

Deleting the Port Security Database

Use the **no port-security** command in configuration mode to delete the configured database for a specified VSAN

```
switch(config)# no port-security database vsan 1
switch(config)#
```

Displaying Port Security Commands

The **show port-security database** commands display the configured port security information (see Examples 18-1 to 18-9). The

Example 18-1 Displays the Contents of the Port Security Database

```
switch# show port-security database
-----
VSAN      Logging-in Entity          Logging-in Point (      Interface)
-----
1         21:00:00:e0:8b:06:d9:1d (pwwn) 20:0d:00:05:30:00:95:de (fc1/13)
1         50:06:04:82:bc:01:c3:84 (pwwn) 20:0c:00:05:30:00:95:de (fc1/12)
2         20:00:00:05:30:00:95:df (swwn) 20:0c:00:05:30:00:95:de (port-channel 128)
3         20:00:00:05:30:00:95:de (swwn) 20:01:00:05:30:00:95:de (fc1/1)
[Total 4 entries]
```

You can optionally specify a fWWN and a VSAN, or an interface and a VSAN in the **show port-security** command to view the output of the activated port security (see Example 18-2).

Example 18-2 Displays the Port Security Database in VSAN 1

```
switch# show port-security database vsan 1
-----
Vsan      Logging-in Entity          Logging-in Point      (Interface)
-----
```

```

1          *                20:85:00:44:22:00:4a:9e (fc3/5)
1      20:11:00:33:11:00:2a:4a (pwwn)  20:81:00:44:22:00:4a:9e (fc3/1)
[Total 2 entries]

```

Example 18-3 Displays the Activated Database

```
switch# show port-security database active
```

```

-----
VSAN      Logging-in Entity                Logging-in Point(   Interface)      Learnt
-----
1         21:00:00:e0:8b:06:d9:1d (pwwn)  20:0d:00:05:30:00:95:de (fc1/13)      Yes
1         50:06:04:82:bc:01:c3:84 (pwwn)  20:0c:00:05:30:00:95:de (fc1/12)      Yes
2         20:00:00:05:30:00:95:df (swwn)  20:0c:00:05:30:00:95:de (port-channel 128) Yes
3         20:00:00:05:30:00:95:de (swwn)  20:01:00:05:30:00:95:de (fc1/1)
[Total 4 entries]

```

The access information for each port can be individually displayed. If you specify the fwwn or interface options, all devices that are paired in the active database (at that point) with the given fwwn or the interface are displayed (see Examples 18-4 to 18-6).

Example 18-4 Displays the Wildcard fWWN Port Security in VSAN 1

```
switch# show port-security database fwwn 20:85:00:44:22:00:4a:9e vsan 1
Any port can login thru' this fwwn
```

Example 18-5 Displays the Configured fWWN Port Security in VSAN 1

```
switch# show port-security database fwwn 20:01:00:05:30:00:95:de vsan 1
20:00:00:0c:88:00:4a:e2 (swwn)
```

Example 18-6 Displays the Interface Port Information in VSAN 2

```
switch# show port-security database interface fc 1/1 vsan 2
20:00:00:0c:88:00:4a:e2 (swwn)
```

The port security statistics are constantly updated and available at any time (see Example 18-7).

Example 18-7 Displays the Port Security Statistics

```

switch# show port-security statistics
Statistics For VSAN: 1
-----
Number of pWWN permit: 2
Number of nWWN permit: 2
Number of sWWN permit: 2
Number of pWWN deny   : 0
Number of nWWN deny   : 0
Number of sWWN deny   : 0

Total Logins permitted : 4
Total Logins denied    : 0
Statistics For VSAN: 2
-----
Number of pWWN permit: 0
Number of nWWN permit: 0
Number of sWWN permit: 2
Number of pWWN deny   : 0
Number of nWWN deny   : 0

```

```
Number of swwn deny : 0
...
```

To verify the status of the active database and the auto-learn configuration, use the **show port-security status** command (see [Example 18-8](#)).

Example 18-8 Displays the Port Security Status

```
switch# show port-security status
VSAN 1 :Activated database, auto-learning is enabled
VSAN 2 :No Active database, auto-learning is disabled
...
```

The **show port-security** command displays the previous 100 violations by default. (see [Example 18-9](#)).

Example 18-9 Displays the Violations in the Port Security Database

```
switch# show port-security violations
```

VSAN	Interface	Logging-in Entity	Last-Time	[Repeat count]
1	fc1/13	21:00:00:e0:8b:06:d9:1d (pwwn)	Jul 9 08:32:20 2003	[20]
		20:00:00:e0:8b:06:d9:1d (nwwn)		
1	fc1/12	50:06:04:82:bc:01:c3:84 (pwwn)	Jul 9 08:32:20 2003	[1]
		50:06:04:82:bc:01:c3:84 (nwwn)		
2	port-channel 1	20:00:00:05:30:00:95:de (swwn)	Jul 9 08:32:40 2003	[1]
[Total 2 entries]				

The **show port-security** command issued with the **last number** option displays only the specified number of entries that appear first.

Default Port Security Settings

[Table 18-4](#) lists the default settings for all security features in any switch.

Table 18-4 Default Security Settings

Parameters	Default
Auto-learn	Enabled if port security is enabled
Port security	Disabled



Configuring Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. Specifically, FSPF is used to:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path.
 - FSPF supports multiple paths.
 - FSPF automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

- [FSPF Features, page 19-2](#)
- [FSPF Examples, page 19-2](#)
- [Configuring FSPF Globally, page 19-4](#)
- [Configuring FSPF for a Specific Interface, page 19-6](#)
- [Configuring Fibre Channel Routes, page 19-8](#)
- [Clearing FSPF Counters, page 19-9](#)
- [Broadcast Routing, page 19-10](#)
- [In-Order Delivery, page 19-10](#)
- [Configuring Flow Statistics, page 19-14](#)
- [Displaying Routing and Forwarding Information, page 19-16](#)
- [Default Settings, page 19-20](#)

FSPF Features

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra's algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra's algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.



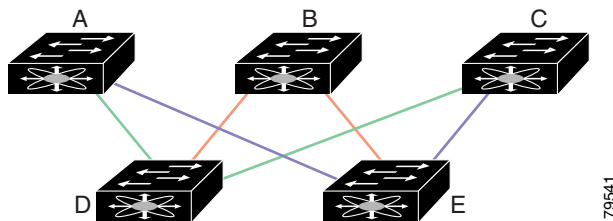
Note

The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 19-1 depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

Figure 19-1 Fault Tolerant Fabric



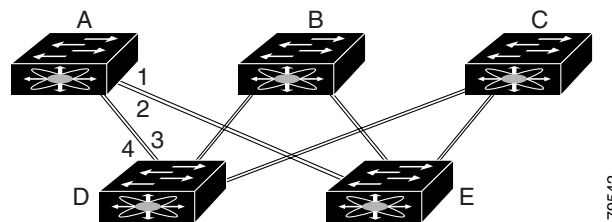
For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in [Figure 19-1](#), each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. [Figure 19-2](#) shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire PortChannel. This configuration also improves the resiliency of the network. The failure of a link in a PortChannel does not trigger a route change, thereby reducing the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 19-2 Fault Tolerant Fabric with Redundant Links



For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in [Figure 19-3](#) and summarized in [Table 19-3](#) and [Table 19-4](#). Two links between switch 1 and switch 2 exist as either equal-cost ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100 utilization of 1G in two scenarios:

- Disabling the traffic link by either physically removing the cable (see [Table 19-3](#)).
- Shutting down either switch 1 or switch 2 (see [Table 19-4](#)).

Figure 19-3 Fail-over Scenario Using Traffic Generators

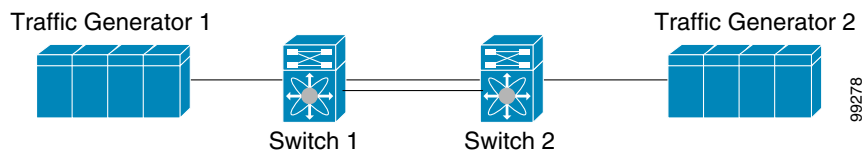


Table 19-1 Physically Removing the Cable for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	Switch 2	Switch 1	Switch 2
110ms (~2K frame drops)		130ms+ (~4k frame drops)	
100 ms hold time when a signal loss is reported as mandated by the standard			

Table 19-2 Shutting Down the Switch for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	Switch 2	Switch 1	Switch 2
~0ms (~8 frame drops)	110ms (~2K frame drops)	130ms+ (~4K frame drops)	
No hold time needed.	Signal loss on switch 1	No hold time needed.	Signal loss on switch 1

Configuring FSPF Globally

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you don't have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.



Note

FSPF is enabled by default. Generally, you do not need to configure these advanced features.



Caution

The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

To configure a FSPF feature for the entire VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# fspf config vsan 1 switch-config- (fspf-config)#	Enters FSPF global configuration mode for the specified VSAN.
Step 3	switch-config- (fspf-config)# spf static switch-config- (fspf-config)#	Forces static SPF computation for the dynamic (default) incremental VSAN.
Step 4	switch-config- (fspf-config)# spf hold-time 10 switch-config- (fspf-config)#	Configures the hold time between two route computations in milliseconds for the entire VSAN. The default value is 0. Note If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly.
Step 5	switch-config- (fspf-config)# region 7 switch-config- (fspf-config)#	Configures the autonomous region for this VSAN and specifies the region ID (7).

Deleting the Entire FSPF Configuration

To delete FSPF configuration for the entire VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# no fspf config vsan 3 switch(config)#	Deletes the FSPF configuration for VSAN 3.

Disabling FSPF Routing Protocols

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

To enable or disable FSPF routing protocols, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# no fspf enable vsan 5 switch(config)#	Disables FSPF routing protocol in VSAN 5.
	switch(config)# fspf enable vsan 7 switch(config)#	Enables FSPF routing protocol in VSAN 7.

Link State Record Defaults

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. [Table 19-3](#) displays the default settings for switch responses.

Table 19-3 LSR Default Settings

LSR Option	Default	Description
Acknowledgement interval (RxmtInterval)	5 seconds	The time a switch waits for an acknowledgement from the LSR before retransmission.
Refresh time (LSRefreshTime)	30 minutes	The time a switch waits before sending an LSR refresh transmission.
Maximum age (MaxAge)	60 minutes	The time a switch waits before dropping the LSR from the database.

Configuring FSPF for a Specific Interface

Several FSPF commands are available on a per interface basis. The following configuration procedures apply to an interface in a specific VSAN and are described in this section.

- [Computing Route Cost, page 19-6](#)
- [Specifying Hello Time Intervals, page 19-6](#)
- [Specifying Dead Intervals, page 19-7](#)
- [Disabling FSPF for Specific Interfaces, page 19-7](#)
- [Retransmitting Intervals, page 19-8](#)

Computing Route Cost

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535 seconds. The default cost for 1Gbps is 1000 and 2Gbps is 500 seconds

To configure FSPF link cost, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc1/4 switch(config-if)#	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# fspf cost 5 vsan 90 switch(config-if)#	Configures the cost for the selected interface in VSAN 90.

Specifying Hello Time Intervals

You can set the FSPF hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.

To configure the FSPF Hello time interval, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc1/4 switch(config-if)#	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# fspf hello-interval 15 vsan 175 switch(config-if)#	Specifies the hello message interval (15 seconds) to verify the health of the link in VSAN 175. The default is 20 seconds.

Specifying Dead Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.



Caution

An error is reported at the command prompt if the configured dead time interval is less than the Hello time interval.

To configure the FSPF dead time interval, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc1/4 switch(config-if)#	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# fspf dead-interval 25 vsan 7 switch(config-if)#	Specifies the maximum interval for VSAN 7 before which a Hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds.

Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.



Note

FSPF must be enabled at both ends of the interface for the protocol to work.

To disable FSPF for a specific interface, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc1/4 switch(config-if)#	Configures a specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# fspf passive vsan 1 switch(config-if)#	Disables the FSPF protocol for the specified interface in the specified VSAN.
	switch(config-if)# no fspf passive vsan 1 switch(config-if)#	Reenables the FSPF protocol for the specified interface in the specified VSAN.

Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.


Note

This value must be the same on the switches on both ends of the interface.

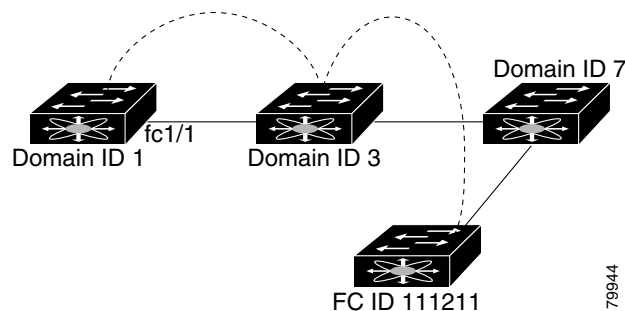
To configure the FSPF retransmit time interval, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc1/4 switch(config-if)#	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# fspf retransmit-interval 15 vsan 12 switch(config-if)#	Specifies the retransmit time interval for unacknowledged link state updates in VSAN 12. The default is 5 seconds.

Configuring Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. To configure the FC ID for the specified interface and domain, you can configure the specified route (for example FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see [Figure 19-4](#)).

Figure 19-4 Fibre Channel Routes


Note

Other than in VSANs, run time checks are not performed on configured and suspended static routes.

To configure an FC route, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# fcroute 0x111211 interface fc1/1 domain 3 vsan 2 switch(config)#	Configures the route for the specified Fibre Channel interface and domain. In this example, interface fc1/1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch.
	switch(config)# fcroute 0x111211 interface port-channel 1 domain 3 vsan 4 switch(config)#	Configures the route for the specified PortChannel interface and domain. In this example, interface port-channel 1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch.
	switch(config)# fcroute 0x031211 interface fc1/1 domain 3 metric 1 vsan 1 switch(config-if)#	Configures the static route for a specific FC ID and next hop domain ID and also assigns the cost of the route. If the remote destination option is not specified, the default is direct.
	switch(config)# fcroute 0x111112 interface fc1/1 domain 3 metric 3 remote vsan 3	Adds a static route to the RIB. If this is an active route and the FIB ¹ records are free, it is also added to the FIB. If the cost (metric) of the route is not specified, the default is 10.
Step 3	switch(config)# fcroute 0x610000 0xff0000 interface fc 1/1 domain 1 vsan 2 switch(config)#	Configures the netmask for the specified route in interface fc1/1 (or PortChannel). You can specify one of three routes: ff0000 matches only the domain, ffff00 matches the domain and the area, ffffff matches the domain, area, and port.

1. FIB = Forwarding Information Base

Clearing FSPF Counters

To clear the FSPF statistics counters for one interface or for the entire VSAN, follow this step:

	Command	Purpose
Step 1	switch# clear fspf counters vsan 1 switch#	Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared.
	switch# clear fspf counters vsan 200 interface fc1/1 switch#	Clears the FSPF statistics counters for the specified interface in VSAN 200.

Broadcast Routing

Broadcast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric (for broadcast traffic).

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive the distribution tree information. The protocols create a loop-free broadcast distribution tree.



Caution

All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

In-Order Delivery

In-order delivery of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

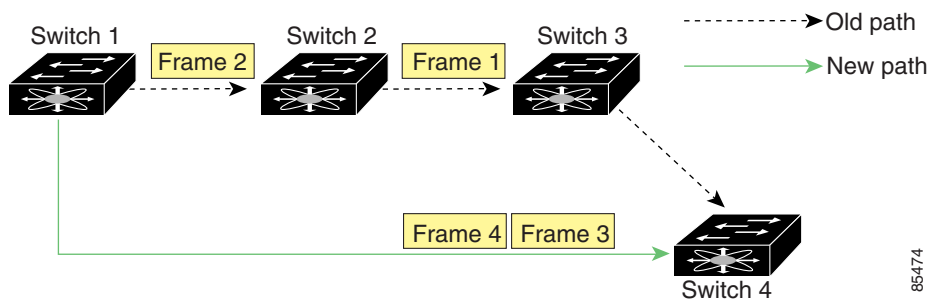
Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

In case of a single switch, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Reordering Network Frames

When you experience a route change in the network. The new selected path may be faster or less congested than the old route (see [Figure 19-5](#)).

Figure 19-5 Route Change Delivery



In [Figure 19-5](#), the new path from Switch 1 to Switch 4 is faster. Hence, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

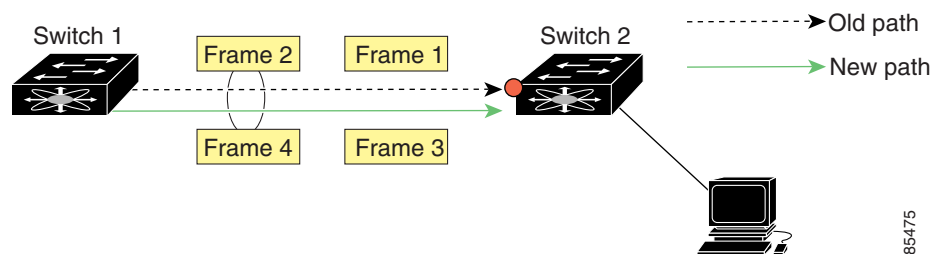
If the in-order guarantee feature is enabled, the frames within the network are treated as specified below:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames which can not be delivered in-order within the network latency drop period are dropped inside the network.
- The number of dropped frames are reduced by slowing down the traffic at the frame source.

Reordering PortChannel Frames

When a link change occurs in a PortChannel, the frames for the same exchange or the same flow can switch from one path to another faster path (see [Figure 19-6](#)).

Figure 19-6 Link Congestion Delivery



In [Figure 19-6](#), the port of the old path (red dot) is congested. Hence Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order guarantee feature is enabled, the frames crossing a PortChannel are treated as specified below:

- Frames using the old path are delivered before new frames are accepted.
- Frames which cannot be delivered in-order, through the old path, within the switch latency drop period are dropped.
- The new frames are delivered through the new path after the switch latency drop period has elapsed.

Enabling In-Order Delivery

By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.



Tip

We recommend that you only enable this feature in a switch when devices are present in the switch that cannot handle any out-of-order frames. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed due to an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out of order.

Use the **show in-order-guarantee** command to display the present configuration status:

```
switch# show in-order-guarantee
```

inorder delivery is not guaranteed

To enable in-order delivery, follow these steps.

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# in-order-guarantee	Enables in-order delivery in the switch.
	switch(config)# no in-order-guarantee	Reverts the switch to the factory defaults and disables the in-order delivery feature.

Configuring the Drop Latency Time

Use this command if you need to change the default latency time for either a network or a switch.

To configure the network and the switch drop latency time, follow these steps.

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcdroplateny network 5000	Configures network drop latency time to be 5000 milliseconds for the network. The valid range is 0 to 60000 milliseconds. The default is 2000 milliseconds. Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network
	switch(config)# no fcdroplateny network 4500	Removes the current fcdroplateny network configuration (4500) and reverts the switch to the factory defaults.
Step 3	switch(config)# fcdroplateny switch 4000	Configures switch drop latency time to be 4000 milliseconds for the switch. The valid range is 0 to 60000 milliseconds. The default is 500 milliseconds. Note The switch drop latency parameter should have the same value in all the switches in the network
	switch(config)# no fcdroplateny switch 4500	Removes the current fcdroplateny switch configuration (4500) and reverts the switch to the factory defaults.

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplateny** command (see [Example 19-1](#)).

Example 19-1 Displays Administrative Distance

```
switch# show fcdroplateny
switch latency value:4000 milliseconds
network latency value:5000 milliseconds
```

Configuring Flow Statistics

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN

If you enable flow counters, you can enable a maximum of 1K entries for aggregate flow and flow statistics. Be sure to assign an unused flow index to a module for each new flow. Flow indexes can be repeated across modules. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

To count the aggregated flow statistics for a VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcflow stats aggregated module 1 index 1005 vsan 1 switch(config)#	Enables the aggregated flow counter.
	switch(config)# no fcflow stats aggregated module 1 index 1005 vsan 1 switch(config)#	Disables the aggregated flow counter.

To count the flow statistics for a source and destination FC ID in a VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcflow stats module 1 index 1 0x145601 0x5601ff ffffff vsan 1 switch(config)#	Enables the flow counter. Note The source ID and the destination ID are specified in the FC ID hex format (for example, 0x123aff). The mask can be one of ff0000 or ffffff.
Step 3	switch(config)# no fcflow stats aggregated module 2 index 1001 vsan 2 switch(config)#	Disables the flow counter.

Clearing FIB¹ Statistics

To clear the aggregated flow counter, use the **clear fcflow stats** command (see Examples 19-2 and 19-3).

Example 19-2 Clears Aggregated Flow Counters

```
switch# clear fcflow stats aggregated module 2 index 1
```

Example 19-3 Clears Flow Counters for Source and Destination FC IDs

```
switch# clear fcflow stats module 2 index 1
```

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics (see Example 19-4 to 19-6).

Example 19-4 Displays Aggregated fcflow Details for the Specified Module

```
switch# show fcflow stats aggregated module 2
Idx  VSAN # frames # bytes
----  ---  -
0000 4      387,653  674,235,875
0001 6      34,402   2,896,628
```

Example 19-5 Displays fcflow Details for the Specified Module

```
switch# show fcflow stats module 2
Idx  VSAN D ID          S ID          mask          # frames # bytes
----  ---  -
0000 4      032.001.002  007.081.012  ff.ff.ff      387,653  674,235,875
0001 6      004.002.001  019.002.004  ff.00.00      34,402   2,896,628
```

Example 19-6 Displays fcflow Index Usage for the Specified Module

```
switch# show fcflow stats usage module 2
2 flows configured
configured flow : 3,7
```

1. FIB = Forwarding Information Base

Displaying Routing and Forwarding Information

You can view specific information about existing Fibre Channel and FSPF configurations at any time from the EXEC mode. The following **show** commands provide further details on existing Fibre Channel paths and routes (see Examples 19-7 to 19-15).



Note

When the number of routes are displayed in the command output, both visible and hidden routes are include in the total number of routes. While the hidden routes are added to the count, they will not be visible.

Example 19-7 Displays Administrative Distance

```
switch# show fcroute distance
```

UUID	Route Distance	Name
----	-----	----
10	20	RIB
22	40	FCDOMAIN
39	80	RIB-CONFIG
12	100	FSPF
17	120	FLOGI
21	140	TLPM
14	180	MCAST
64	200	RIB-TEST

Example 19-8 Displays Multicast Routing Information

```
switch# show fcroute multicast
```

VSAN	FC ID	# Interfaces
----	-----	-----
1	0xffffffff	0
2	0xffffffff	1
3	0xffffffff	1
4	0xffffffff	0
5	0xffffffff	0
6	0xffffffff	0
7	0xffffffff	0
8	0xffffffff	0
9	0xffffffff	0
10	0xffffffff	0

Example 19-9 Displays FCID Information for a Specified VSAN

```
switch# show fcroute multicast vsan 3
```

VSAN	FC ID	# Interfaces
----	-----	-----
3	0xffffffff	1

Example 19-10 Displays FCID and interface Information for a Specified VSAN

```
switch# show fcroute multicast 0xffffffff vsan 2
```

VSAN	FC ID	# Interfaces
----	-----	-----
2	0xffffffff	1
	fc1/1	

Example 19-11 Displays Unicast Routing Information

```
switch# show fcroute unicast
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops   Cost
-----
static  1    0x010101 0xffffffff 0x00 0x00 D P A 1    10
static  2    0x111211 0xffffffff 0x00 0x00 R P A 1    10
fspf    3    0x610000 0xff0000 0x00 0x00 D P A 4    500
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040102 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040103 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040104 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x111211 0xffffffff 0x00 0x00 D P A 1    10
```

Example 19-12 Displays Unicast Routing Information for a Specified VSAN

```
switch# show fcroute unicast vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops   Cost
-----
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040102 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040103 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040104 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x111211 0xffffffff 0x00 0x00 D P A 1    10
```

Example 19-13 Displays Unicast Routing Information for a Specified FCID

```
switch# show fcroute unicast 0x040101 0xffffffff vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops   Cost
-----
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
    fc1/2 Domain 0xa6(166)
```

Example 19-14 Displays Route Database Information

```
switch# show fcroute summary
FC Route Database Created Thu Feb 13 07:21:52 2003
VSAN      Ucast      Mcast      Label      Last Modified Time
-----
1          5          1          0          Thu Feb 13 10:21:06 2003
2          4          1          0          Thu Feb 13 10:21:07 2003
3          4          1          0          Thu Feb 13 10:21:08 2003
4          4          1          0          Thu Feb 13 10:21:09 2003
5          4          1          0          Thu Feb 13 10:21:10 2003
6          4          1          0          Thu Feb 13 10:21:11 2003
7          4          1          0          Thu Feb 13 10:21:12 2003
8          4          1          0          Thu Feb 13 10:21:13 2003
9          4          1          0          Thu Feb 13 10:21:14 2003
10         4          1          0          Thu Feb 13 10:21:15 2003
11         4          1          0          Thu Feb 13 10:21:16 2003
12         4          1          0          Thu Feb 13 10:21:17 2003
13         4          1          0          Thu Feb 13 10:21:18 2003
14         4          1          0          Thu Feb 13 10:21:18 2003
15         4          1          0          Thu Feb 13 10:21:19 2003
-----
Total      61          15          0
```

Example 19-15 Displays Route Database Information for a Specified VSAN

```
switch# show fcroute summary vsan 5
FC Route Database Created Thu Feb 13 07:21:52 2003
```

VSAN	Ucast	Mcast	Label	Last Modified Time
5	4	1	0	Thu Feb 13 10:21:10 2003
Total	4	1	0	

Displaying Global FSPF Information

The **show fspf** command (see [Example 19-16](#)) displays global FSPF information for a specific VSAN:

- the domain number of the switch
- the autonomous region for the switch
- Min_LS_arrival: the minimum time that must elapse before the switch accepts LSR updates
- Min_LS_interval: the minimum time that must elapse before the switch can transmit an LSR
- LS_refresh_time: the interval lapse between refresh LSR transmissions
- Max_age: the maximum time aa LSR can stay before being deleted

Example 19-16 Displays FSPF Information for a Specified VSAN

```
switch# show fspf vsan 1
FSPF routing for VSAN 1
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 5000 msec
Local Domain is 0x65(101)
Number of LSRs = 3, Total Checksum = 0x0001288b

Protocol constants :
  LS_REFRESH_TIME = 1800 sec
  MAX_AGE          = 3600 sec

Statistics counters :
  Number of LSR that reached MaxAge = 0
  Number of SPF computations         = 7
  Number of Checksum Errors          = 0
  Number of Transmitted packets :   LSU 65 LSA 55 Hello 474 Retranmsitted LSU 0
  Number of received packets :     LSU 55 LSA 60 Hello 464 Error packets 10
```

Displaying the FSPF Database

The **show fspf database** command displays a summary of the FSPF database for a specified VSAN (see [Example 19-17](#)). If other parameters are not specified, all LSRs in the database are displayed:

- LSR Type
- Domain ID of the LSR owner
- Domain ID of the advertising router

- LSR age
- LS incarnation member
- Number of links

You could narrow the display to obtain specific information by issuing additional parameters for the domain ID of the LSR owner. For each interface, the following information is also available:

- Domain ID of the neighboring switch
- E port index
- Port index of the neighboring switch
- Link type and cost

Example 19-17 Displays FSPF Database Information

```
switch# show fspf database vsan 1
```

```
FSPF Link State Database for VSAN 1 Domain 0x0c(12)
```

```
LSR Type           = 1
Advertising domain ID = 0x0c(12)
LSR Age            = 1686
LSR Incarnation number = 0x80000024
LSR Checksum       = 0x3caf
Number of links     = 2
```

NbrDomainId	IfIndex	NbrIfIndex	Link Type	Cost
0x65(101)	0x0000100e	0x00001081	1	500
0x65(101)	0x0000100f	0x00001080	1	500

```
FSPF Link State Database for VSAN 1 Domain 0x65(101)
```

```
LSR Type           = 1
Advertising domain ID = 0x65(101)
LSR Age            = 1685
LSR Incarnation number = 0x80000028
LSR Checksum       = 0x8443
Number of links     = 6
```

NbrDomainId	IfIndex	NbrIfIndex	Link Type	Cost
0xc3(195)	0x00001085	0x00001095	1	500
0xc3(195)	0x00001086	0x00001096	1	500
0xc3(195)	0x00001087	0x00001097	1	500
0xc3(195)	0x00001084	0x00001094	1	500
0x0c(12)	0x00001081	0x0000100e	1	500
0x0c(12)	0x00001080	0x0000100f	1	500

```
FSPF Link State Database for VSAN 1 Domain 0xc3(195)
```

```
LSR Type           = 1
Advertising domain ID = 0xc3(195)
LSR Age            = 1686
LSR Incarnation number = 0x80000033
LSR Checksum       = 0x6799
Number of links     = 4
```

NbrDomainId	IfIndex	NbrIfIndex	Link Type	Cost
0x65(101)	0x00001095	0x00001085	1	500
0x65(101)	0x00001096	0x00001086	1	500
0x65(101)	0x00001097	0x00001087	1	500
0x65(101)	0x00001094	0x00001084	1	500

Displaying FSPF Interfaces

The **show fspf** command displays the following information for each selected interface (see [Example 19-18](#)).

- link cost
- timer values
- neighbor's domain ID (if known)
- local interface number
- remote interface number (if known)
- FSPF state of the interface
- interface counters

Example 19-18 Displays FSPF Interface Information

```
switch# show fspf vsan 1 interface fc1/1
FSPF interface fc1/1 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x0c(12), Neighbor Interface index is 0x0f100000
Statistics counters :
    Number of packets received : LSU 8 LSA 8 Hello 118 Error packets 0
    Number of packets transmitted : LSU 8 LSA 8 Hello 119 Retransmitted LSU 0
    Number of times inactivity timer expired for the interface = 0
```

Default Settings

[Table 19-4](#) lists the default settings for FSPF features.

Table 19-4 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgement interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).
Routing table	FSPF stores up to 16 equal cost paths to a given destination.

Table 19-4 Default FSPF Settings (continued)

Parameters	Default
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.



Configuring IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding or in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

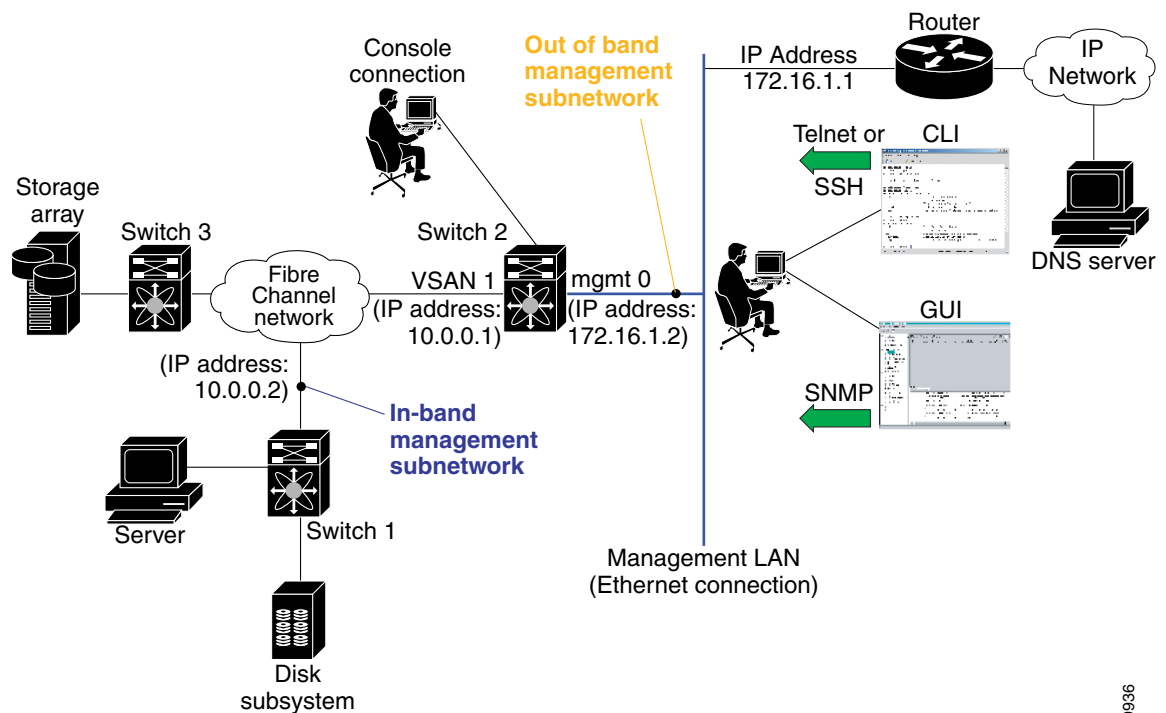
This chapter includes the following sections:

- [Traffic Management Services, page 20-2](#)
- [Configuring the Ethernet Management Port, page 20-2](#)
- [Configuring the Default Gateway, page 20-3](#)
- [Configuring the Default Network, page 20-4](#)
- [IP Access Control Lists, page 20-5](#)
- [Configuring IPFC, page 20-11](#)
- [Configuring IP Static Routes, page 20-12](#)
- [Displaying IP Interface Information, page 20-13](#)
- [Configuring Overlay VSANs, page 20-14](#)
- [Configuring Multiple VSANs, page 20-16](#)
- [Configuring VRRP, page 20-18](#)
- [Configuring DNS Server, page 20-24](#)
- [Default Settings, page 20-25](#)

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running IP protocol over a FC interface can access the switch using the IPFC functionality. If the NMS does not have a Fibre Channel HBA, in-band management can still be performed using one of the switches as an access point to the fabric (see Figure 20-1).

Figure 20-1 Management Access to Switches



79936

Configuring the Ethernet Management Port

The management port on the switch allows multiple simultaneous Telnet or SNMP network management sessions. You can also configure the supervisor module's Ethernet interface and VSAN interfaces as management ports. This section focuses on the Ethernet management port (mgmt0). You can remotely configure the switch through the management port. To configure a connection remotely, you must configure the IP parameters (IP address and subnet mask) from the CLI so that the switch is reachable.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

To configure the mgmt0 Ethernet interface, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface mgmt0 switch(config-if)#	Enters the interface configuration mode on the management Ethernet interface (mgmt0).
Step 3	switch(config-if)# ip address 1.1.1.1 255.255.255.0	Enters the IP address (1.1.1.1) and IP subnet mask (255.255.255.0) for the management interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

Configuring the Default Gateway

Use the **IP default-gateway** command to configure the IP address for a switch's default gateway. This IP address should be configured along with the IP static routing commands (IP default-network, destination prefix, and destination mask, and next hop address)

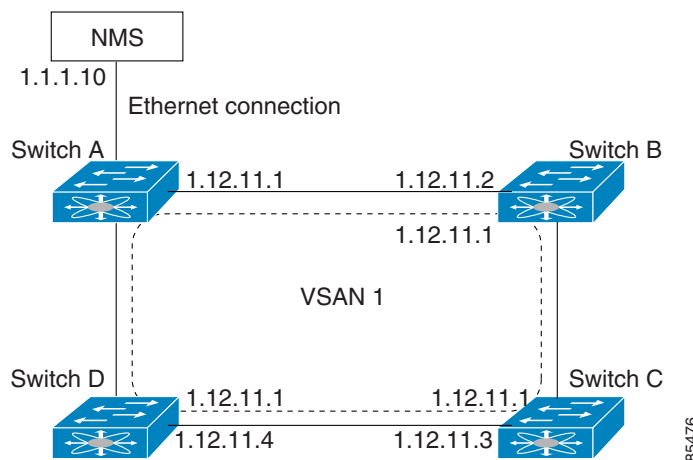


Tip

If you configure the static route IP forwarding and the default-network details, these IP addresses will be used regardless of the default-gateway being enabled or disabled. If these IP address are configured and not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch (see the [“Initial Setup Routine”](#) section on page 4-2).

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IP address of the gateway switch (see [Figure 20-2](#)).

Figure 20-2 Overlay VSAN Functionality



In [Figure 20-2](#), switch A has the IP address 1.12.11.1, switch B has the IP address 1.12.11.2, switch C has the IP address 1.12.11.3, and switch D has the IP address 1.12.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IP address, 1.12.11.1, in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet world, the frame is routed through the gateway switch.

When forwarding is disabled (default), IP frames are not sent from one interface to another. In these cases, the software performs local IP routing between two switches using the in-band option for Fibre Channel traffic and the mgmt0 option for Ethernet traffic.

When a VSAN is created, a VSAN interface is not created automatically. You need to specifically create the interface (see the [“Configuring VSAN Interfaces”](#) section on page 10-17).

To configure default gateways, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip default-gateway 1.12.11.1 switch(config)#	Configures the IP address for the default gateway (1.12.11.1).

Use the **show ip route** command to verify that the IP address for the default gateway is configured.

Configuring the Default Network

Unlike the **ip default-gateway** command, use the **ip default-network** command when IP routing is enabled on the switch. If you assign the IP default network address, the switch considers routes to that network as the last resort. If the IP default network address is not available, the switch uses the IP default gateway address. For every network configured with the IP default network address, the switch flags that route as a candidate default route, if the route is available.



Tip

If you configure the static route IP forwarding and the default network details, these IP addresses will be used regardless of the default gateway being enabled or disabled. If these IP address are configured and not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch (see the [“Initial Setup Routine”](#) section on page 4-2).

To configure default networks, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip default-network 190.10.1.0 switch(config)#	Configures the IP address for the default network (190.10.1.0).
	switch(config)# ip route 10.0.0.0 255.0.0.0 131.108.3.4 switch(config)# ip default-network 10.0.0.0 switch(config)#	Defines a static route to network 10.0.0.0 as the static default route.

Use the **show ip route** command to verify if the IP address for the default gateway is configured.

IP Access Control Lists

IP Access control lists (IP-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related MDS out-of-band management traffic and in-band traffic based on IP addresses (Layer 3 and Layer 4 information).

You can use IP-ACLs to control transmissions on an interface.

IP-ACL Configuration Guidelines

Follow these guidelines when configuring IP-ACLs in any switch or director in the Cisco MDS 9000 Family:

- IP-ACLs cannot be configured for Gigabit Ethernet interfaces (IPS modules) or for Fibre Channel interfaces.
- IP-ACLs can only be configured on the management interface and VSAN interfaces.
- An IP-ACL is a sequential collection of permit and deny conditions that apply to IP flows. Each IP packet is tested against the conditions in the list. The first match determines if the software accepts or rejects the rule. Because the software stops testing conditions after the first match, the order of the conditions in the list is critical. If no conditions match, the software rejects that rule.
- An IP protocol can be configured using an integer ranging from 0 to 255 to represent a particular IP protocol. Alternatively, you can specify the name of a protocol: **icmp**, **ip**, **tcp**, or **udp**. IP includes Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), and other protocols.
- The source/source-wildcard and destination/destination-wildcard is specified in one of two ways:
 - Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
 - Using the **any** option as an abbreviation for a source/source-wildcard or destination/destination-wildcard (0.0.0.0/255.255.255.255)
- To configure an IP-ACL, you must complete the following tasks:
 1. Create an IP-ACL by specifying a name and access condition.

All lists use the source and destination address for matching operations. You can configure finer granularity using optional keywords
 2. Apply the access list to specified interfaces.

Creating IP-ACLs

You can specify IP-ACLs using a assigned name. Each IP-ACL can have a maximum of 256 entries. Each entry is a unique filter applied to a specified interface. Each switch can have a maximum of 64 IP-ACLs.

Traffic coming into the switch is compared to IP-ACL entries based on the order that the entries occur in the switch. New statements are added to the end of the list. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the list, the traffic is denied. For this reason, you should have the frequently hit entries at the top of the list. There is an *implied deny* for traffic that is not permitted. A single-entry IP-ACL with only one **deny** entry has the effect of denying all traffic.

To create an IP-ACL, follow these steps:

	Command	Purpose
Step 1	switch# <code>config t</code>	Enters configuration mode.
Step 2	switch(config)# <code>ip access-list List1 permit ip any any</code>	Configures an IP-ACL called List1 and permits IP traffic from any source address to any destination address.
	switch(config)# <code>no ip access-list List1 permit ip any any</code>	Removes the IP-ACL called List1.
Step 3	switch(config)# <code>ip access-list List1 deny tcp any any</code>	Updates List1 to deny TCP traffic from any source address to any destination address.

To define an IP-ACL that permits a specified network, follow these steps:

	Command	Purpose
Step 1	switch# <code>config t</code>	Enters configuration mode.
Step 2	switch(config)# <code>ip access-list List1 permit udp 192.168.32.0 0.0.7.255</code>	Defines an IP-ACL that permits this network. Subtracting 255.255.248.0 (normal mask) from 255.255.255.255 yields 0.0.7.255.

Adding Entries to an Existing IP-ACL

After you create an IP-ACL, you place subsequent additions at the end of the IP-ACL. You cannot insert entries in the middle of an IP-ACL. Each configured entry is automatically added to the end of a IP-ACL.

To add entries to an existing IP-ACL, follow these steps:

	Command	Purpose
Step 1	switch# <code>config t</code>	Enters configuration mode.
Step 2	switch(config)# <code>ip access-list List1 permit tcp host 10.1.1.2 host 172.16.1.1 eq telnet</code> switch(config)# <code>ip access-list List1 permit tcp host 10.1.1.2 host 172.16.1.1</code> switch(config)# <code>ip access-list List1 permit udp host 10.1.1.2 host 172.16.1.1</code> switch(config)# <code>ip access-list List1 permit ip 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255</code>	Permits all IP traffic from and to the specified networks. Note In this example, the last entry is sufficient. You do not need the first three entries.

Comparing Ports

Use the following operators to compare the source and destination ports:

- **eq** = equal
- **gt** = greater than
- **lt** = less than
- **range** = range of ports

To use the operand and port options, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any	Denies TCP traffic from 1.2.3.0 through source port 5 to any destination.

Port numbers range from 0 to 65535 for TCP and UDP ports. [Table 20-1](#) displays the port numbers for associated TCP and UDP ports.

Table 20-1 TCP and UDP Port Numbers

Protocol	Port	Number
Note If the TCP connection is already established, use the established option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	radius	1812
	wbem-http	5988
	wbem-https	5989
UDP	dns	53
	tftp	69
	ntp	123
	snmp	161
	snmp-trap	162
	syslog	514

ICMP packets are filtered by the ICMP message type or the message code. Both values range from 0 to 255. [Table 20-2](#) displays the value for each associated ICMP type.

Table 20-2 ICMP Type Value

ICMP Type ¹	Value
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

1. ICMP redirect packets are always rejected.

Removing Entries from an Existing IP-ACL

Use the **no permit** and **no deny** commands to remove entries from a configured IP-ACL.

To remove configured entries from an IP-ACL, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# no ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any	Removes this entry from the IP-ACL.
	switch(config)# no ip access-list x3 deny ip any any	Removes this entry from the IP-ACL.
	switch(config)# no ip access-list x3 permit ip any any	Removes this entry from the IP-ACL.

Applying IP-ACLs

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to the switch's interface.

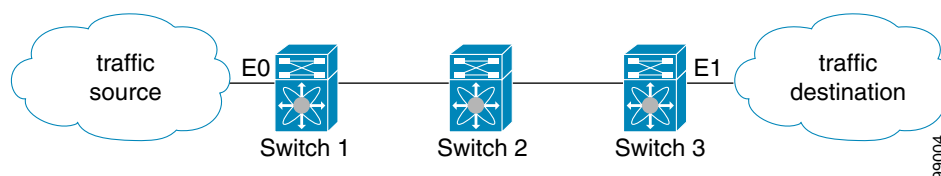


Tip

Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IP-ACL to E0 on Switch 1 instead of an outbound list to E1 on Switch 3 (see [Figure 20-1](#)).

Figure 20-3 Denying Traffic on the Inbound Interface



The **access-group** command controls access to an interface. Each interface can only be associated with one access list per direction. The ingress direction can have a different ACL than the egress direction. The access group becomes active on creation.

**Tip**

We recommend creating all rules in an access list, before creating the access group that uses this access-list.

**Caution**

If you create an access group before an access-list, all packets in that interface are dropped, because the access list is empty.

The terms *in*, *out*, *source*, and *destination* are used as referenced by the switch.

- **In**—Traffic that is arriving on the interface and which will go through the switch; the source would be where it's been and the destination is where it's going (on the other side of the router).

**Tip**

The access-group configuration for the ingress traffic applies to both local and remote traffic.

- **Out**—Traffic that has already been through the switch and is leaving the interface; the source would be where it's been (on the other side of the router) and the destination is where it's going.

**Tip**

The access-group configuration for the egress traffic applies only to local traffic.

To create an access group, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface mgmt0 switch(config-if)#	Configures a management interface (mgmt0).
Step 3	switch(config-if)# ip access-group SampleName	Creates an access group called SampleName for both the ingress and egress traffic (default)
	switch(config-if)# no ip access-group NotRequired	Deletes the access group called NotRequired.
Step 4	switch(config-if)# ip access-group SampleName1 in	Creates an access group called SampleName (if it does not already exist) for ingress traffic.
	switch(config-if)# no ip access-group SampleName1 in	Deletes the access group called SampleName for ingress traffic.
	switch(config-if)# ip access-group SampleName2 out	Creates an access group called SampleName (if it does not already exist) for local egress traffic.
	switch(config-if)# no ip access-group SampleName2 out	Deletes the access group called SampleName for local egress traffic.

Using the **log-deny** option at the end of the individual ACL entries shows the ACL number and whether the packet was permitted or denied, in addition to port-specific information. This option causes an information logging message about the packet that matches the dropped entry (or entries).

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not dumped to the log.

Below is an example of an input ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

Below is an example of an output ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00
TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Displaying IP-ACLs

Use the **show ip access-list** command to view the contents of configured access lists. Each access list can have several filters.

Example 20-1 Displays Configured IP-ACLs

```
switch# show ip access-list usage
Access List Name/Number      Filters IF   Status      Creation Time
-----
abc                          3          7    active    Tue Jun 24 17:51:40 2003
x1                           3          1    active    Tue Jun 24 18:32:25 2003
x3                           0          1  not-ready  Tue Jun 24 18:32:28 2003
```

Example 20-2 Displays a Summary of the Specified IP-ACL

```
switch# show ip access-list abc
ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)
```

Clearing IP-ACL Counters

Use the **clear** command to clear the counters for a specified IP-ACL entry. Note that you cannot use this command to clear the counters for each individual filter.

```
switch# clear ip access-list counters abc permit ip 10.1.1.0 0.0.0.255
```


Configuring IPFC

Once the VSAN interface is created, you can specify the IP address for that VSAN using the **ip address** command.

Configuring an IP Address in a VSAN

To configure a VSAN interface and an IP address for that interface, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 1 switch(config-if)#	Configures the interface for the specified VSAN (1).
Step 3	switch(config-if)# ip address 10.0.0.12 255.255.255.0 switch(config-if)#	Configures the IP address and netmask for the selected interface.

Enabling IP Routing

By default, the IP routing feature is disabled in all switches. To enable the IP routing feature, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# ip routing switch(config)#	Enables IP routing (disabled by default).
Step 3	switch(config)# no ip routing switch(config)#	Disables IP routing and reverts to the factory settings.

Configuring IP Static Routes

Static routing is a mechanism to configure IP routes on the switch. You can configure more than one static route.

To configure a static route, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# IP route <network IP address> <netmask> <next hop IP address> distance <number> interface <vsan number> For example: switch(config)# IP route 10.0.0.0 255.0.0.0 20.20.20.10 distance 10 interface vsan 1 switch(config)#	Configures the static route for the specified IP address, subnet mask, next hop, and distance, and VSAN or management interface.

If your configuration does not need an external router, you can use static routing.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IP routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

Viewing and Clearing ARPs

Address Resolution Protocol (ARP) entries can be viewed (**show arp**), deleted (**no arp**), or cleared (**clear arp-cache**) in Cisco MDS 9000 Family switches. The ARP feature is enabled on all switches.

- Use the **show arp** command to display the ARP table.

```
switch# show arp
Protocol Address           Age (min)  Hardware Addr  Type   Interface
Internet 171.1.1.1                0           0006.5bec.699c  ARPA   mgmt0
Internet 172.2.0.1                4           0000.0c07.ac01  ARPA   mgmt0
```

- Use the **no arp** command in configuration mode to remove an ARP entry from the ARP table.

```
switch(config)# no arp 172.2.0.1
switch(config)#
```

- Use the **clear arp** command to delete all entries from the ARP table. The ARP table is empty by default.

```
switch# clear arp-cache
switch#
```

Displaying IP Interface Information

Use the following **show** commands to view configured IP interface information (see Examples 20-3 to 20-6).

Example 20-3 Displays the VSAN Interface

```
switch# show interface vsan1
vsan1 is up, line protocol is up
  WWPN is 10:00:00:05:30:00:59:1f, FCID is 0x9c0100
  Internet address is 10.1.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  0 packets input, 0 bytes, 0 errors, 0 multicast
  0 packets output, 0 bytes, 0 errors, 0 dropped
```



Note You can see the output for this command only if you have previously configured a virtual network interface (see the “Configuring an IP Address in a VSAN” section on page 20-11).

Example 20-4 Displays the Connected and Static Route Details

```
switch# show ip route

Codes: C - connected, S - static

Default gateway is 172.22.95.1

C 172.22.95.0/24 is directly connected, mgmt0
C 10.1.1.0/24 is directly connected, vsan1
```

Example 20-5 Displays Configured Routes

```
switch# show ip route configured
```

Destination	Gateway	Mask	Metric	Interface
default	172.22.95.1	0.0.0.0	0	mgmt0
10.1.1.0	0.0.0.0	255.255.255.0	0	vsan1
172.22.95.0	0.0.0.0	255.255.255.0	0	mgmt0

Example 20-6 Displays the IP Routing Status

```
switch# show ip routing
ip routing is disabled
```

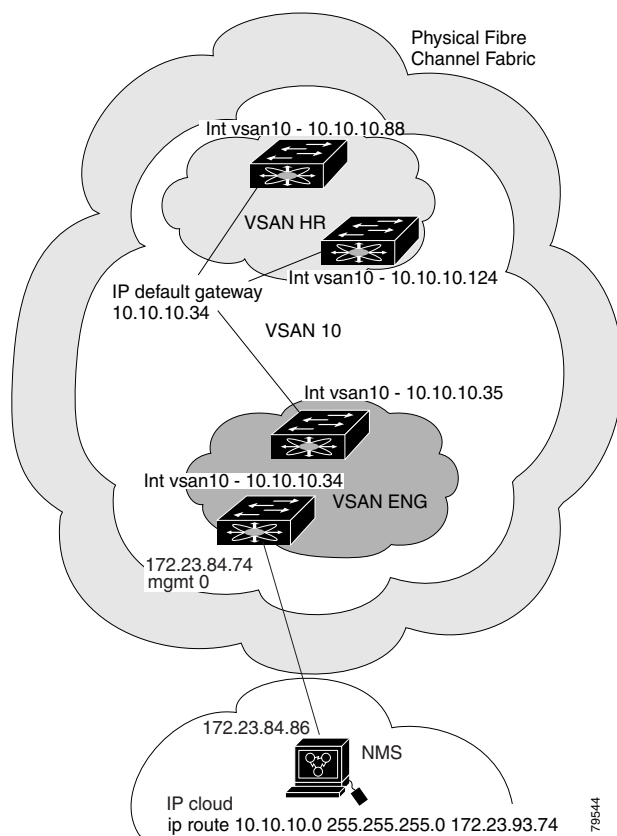
Configuring Overlay VSANs

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.

To configure an overlay VSAN, follow these steps:

- Step 1** Add the VSAN to the VSAN database on all switch in the fabric.
- Step 2** Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side
- Step 3** Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.
- Step 4** Configure default gateway (route) and the IP address on switches that point to the NMS (see [Figure 20-4](#)).

Figure 20-4 Overlay VSAN Configuration Example



The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

To configure an overlay VSAN in one switch (using the example in [Figure 20-4](#)), follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# vsan database switch-config-vsan-db#	Configures the VSAN database.
Step 3	switch--config-vsan-db# vsan 10 name MGMT_VSAN switch--config-vsan-db#	Defines the VSAN in the VSAN database on all of the switches in the Fibre Channel fabric.
Step 4	switch--config-vsan-db# exit switch(config)#	Exits the VSAN database mode.
Step 5	switch(config)# interface vsan10 switch(config-if)#	Creates a VSAN interface (VSAN 10).
Step 6	switch(config-if)# ip address 10.10.10.x netmask 255.255.255.0 switch(config-if)#	Assigns an IP address and netmask on all switches in the fabric.
Step 7	switch(config-if)# no shut	Enables the configured interface.
Step 8	switch--config-if# end switch#	Exits to EXEC mode.
Step 9	switch# exit	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.

To configure the NMS station displayed in [Figure 20-4](#), follow this step:

	Command	Purpose
Step 1	switch# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.



Note

To configure the management interface displayed in [Figure 20-4](#), set the default gateway to an IP address on the Ethernet network.

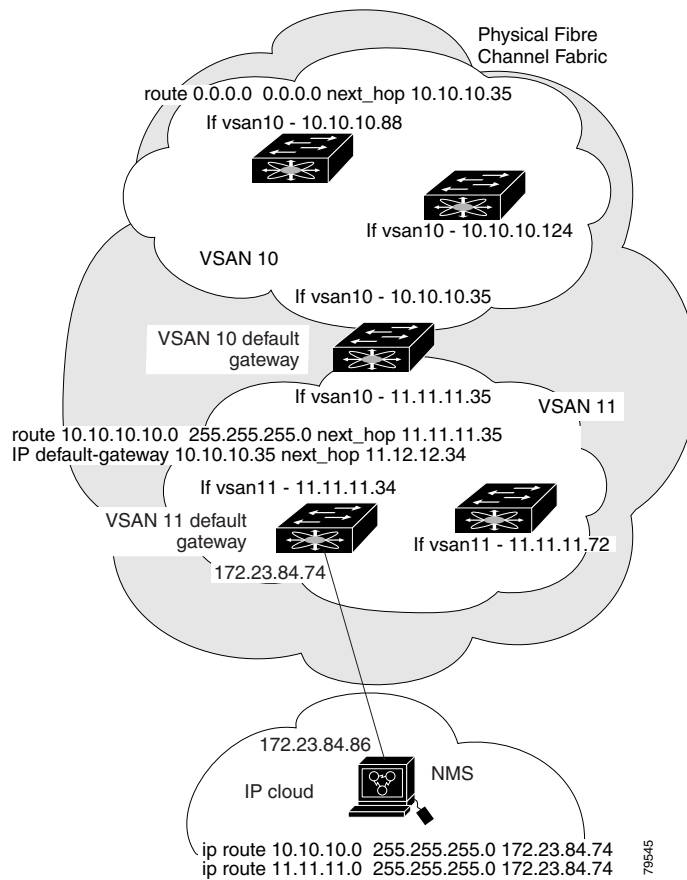
Configuring Multiple VSANs

More than one VSAN can be used to segment the management network in multiple subnets. An active interface must be present on the switch for the VSAN interface to be enabled.

To configure an overlay VSAN, follow these steps:

- Step 1** Add the VSAN to the VSAN database on any switch in the fabric.
- Step 2** Create a VSAN interface for the appropriate VSAN on any switch in the fabric.
- Step 3** Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.
- Step 4** Define the multiple static route on the Fibre Channel switches and the IP cloud (see [Figure 20-5](#)).

Figure 20-5 Multiple VSANs Configuration Example



To configure an overlay VSAN (using the example in [Figure 20-5](#)), follow these steps:

	Command	Purpose
Step 1	switch# confi t	Enters configuration mode.
Step 2	switch(config)# vsan database switch-config-vsan-db#	Configures the VSAN database.

	Command	Purpose
Step 3	switch-config-vsan-db# vsan 10 name MGMT_VSAN_10 switch-config-vsan-db#	Defines the VSAN in the VSAN database on all of the switches in VSAN 10.
Step 4	switch-config-vsan-db# exit switch(config)#	Exits the database 10 mode.
Step 5	switch-config-vsan-db# vsan 11 name MGMT_VSAN_11 switch-config-vsan-db#	Defines the VSAN in the VSAN database on all of the switches in VSAN 11.
Step 6	switch-config-vsan-db# exit switch(config)#	Exits the VSAN database 11 mode.
Step 7	switch(config)# interface vsan10 switch(config-if)#	Enters the VSAN 10 interface configuration mode for VSAN 10.
Step 8	switch(config-if)# ip address 10.10.10.x netmask 255.255.255.0 switch(config-if)#	Assigns an IP address and netmask on all switches in VSAN 10.
Step 9	switch(config-if)# no shut	Enables the configured interface for VSAN 10.
Step 10	switch--config-if# exit switch(config)#	Exits the VSAN 10 interface mode.
Step 11	switch(config)# interface vsan11 switch(config-if)#	Enters the VSAN 11 interface configuration mode.
Step 12	switch(config-if)# ip address 11.11.11.x netmask 255.255.255.0 switch(config-if)#	Assigns an IP address and netmask on all of the switches in VSAN 11.
Step 13	switch(config-if)# no shut	Enables the configured interface for VSAN 11.
Step 14	switch--config-if# end switch#	Exits to EXEC mode.
Step 15	switch# exit	Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric.
Step 16	NMS# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74	Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the IP cloud.
Step 17	NMS# route ADD 11.11.11.0 MASK 255.255.255.0 172.22.93.74	Defines a static route for VSAN 11 on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric.
Step 18	switch# route 10.10.10.0 255.255.255.0 next_hop 11.11.11.35	Defines the route to reach subnet 10 from subnet 11.

Configuring VRRP

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. This section provides details on the VRRP feature.

VRRP Features

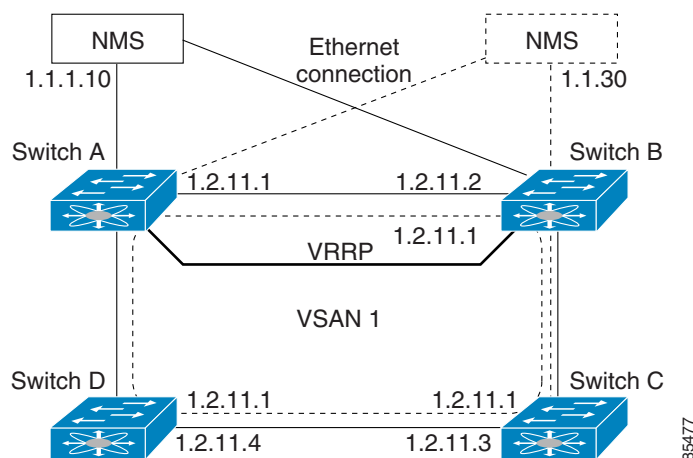
VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following characteristics and advantages:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with a different virtual router IP mapping.
- Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.

VRRP Functionality

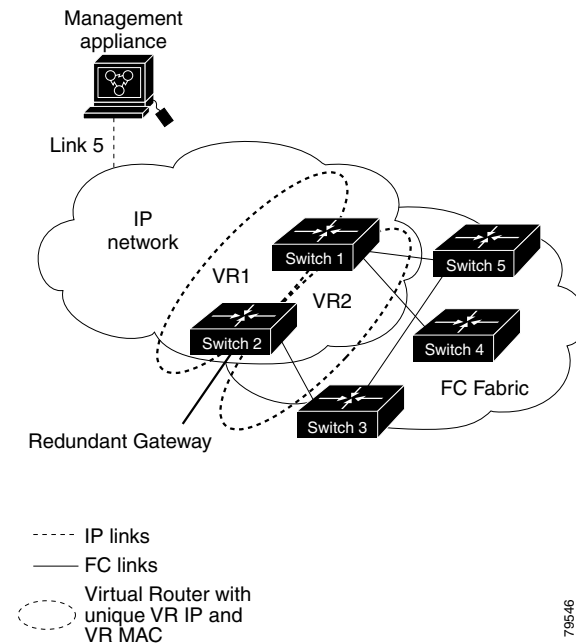
In [Figure 20-6](#), switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches don't have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

Figure 20-6 VRRP Functionality



In [Figure 20-7](#), the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 20-7 Redundant Gateway



Creating or Removing a Virtual Router

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.

To create or remove a VR, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface vsan 1 switch(config-if)#	Configures a VSAN interface (VSAN 1).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)	Creates a VR ID 250.
	switch(config-if-vrrp)# no vrrp 250 switch(config-if)	Removes a VR ID 250.

Enabling a Virtual Router

By default, a virtual router is always disabled (**shutdown**). VRRP can be configured only if this state is disabled. Be sure to configure at least one IP address before attempting to enable a VR.

To enable or disable a virtual router, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if-vrrp)# no shutdown</code>	Enables VRRP configuration.
	<code>switch(config-if-vrrp)# shutdown</code>	Disables VRRP configuration.

Adding an IP Address for a Virtual Router

One primary IP address and multiple secondary addresses can be configured for a switch. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address.

To configure an IP address for a virtual router, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code>	Enters configuration mode.
Step 2	<code>switch(config)# interface vsan 1</code> <code>switch(config-if)#</code>	Configures a VSAN interface (VSAN 1).
Step 3	<code>switch(config-if)# interface ip address 10.0.0.12 255.255.255.0xi</code>	Configures an IP address. The IP address must be configured before the VRRP is added.
Step 4	<code>switch(config-if)# vrrp 250</code> <code>switch(config-if-vrrp)#</code>	Creates VR ID 250.
Step 5	<code>switch(config-if-vrrp)# address 10.0.0.10</code>	Configures the IP address (10.0.0.10) for the selected VR. Note This IP address should be in the same subnet as the IP address of the interface.
	<code>switch(config-if-vrrp)# no address 10.0.0.10</code>	Removes the IP address (10.0.0.10) for the selected VR.

Setting Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for a switch with the primary IP address.

To set the priority for a virtual router, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code>	Enters configuration mode.
Step 2	<code>switch(config)# interface vsan 1</code> <code>switch(config-if)#</code>	Configures a VSAN interface (VSAN 1).

	Command	Purpose
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# priority 2 switch(config-if-vrrp)#	Configures the priority for the selected VRRP. Note Priority 255 cannot be preempted.

Setting the Time Interval for the Advertisement Packet

The valid time range for an advertisement packet is between 1 and 255 seconds with the default being 1 (one) second. If the switch has the primary IP address, this time must be specified.

To set the priority for a virtual router, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 1 switch(config-if)#	Configures a VSAN interface (VSAN 1).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# advertisement-interval 15	Sets the interval time in seconds between sending advertisement frames.

Preempting the Master Virtual Router

By default, the preempt option is enabled. An owner with priority 255 cannot be preempted. If two priorities match, the owner with the highest priority preempts the master virtual router.

To enable or disable preempting, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 1 switch(config-if)#	Configures a VSAN interface (VSAN 1).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# preempt	Enables the higher priority backup virtual router to preempt the lower priority master virtual router. Note This preemption does not apply to the primary IP address.
	switch(config-if-vrrp)# no preempt	Disables the preempt option and allows the master to keep its priority level.



Note

The VRRP **preempt** option is not supported on IP storage Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Configuring Authentication for the Virtual Router

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.

**Note**

All VRRP configurations must be duplicated

To set an authentication option for a virtual router, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 1 switch(config-if)#	Configures a VSAN interface (VSAN 1).
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# authentication text password	Assigns the simple text authentication option and specifies the password for this option.
	switch(config-if-vrrp)# authentication md5 password2003 spi 0x2003	Assigns MD5 authentication option and specifies the key and the unique SPI value for this option. The SPI and the valid range is 0x100 to 0xFFFFFFFF.
	switch(config-if-vrrp)# no authentication	Assigns the no authentication option, which is the default.

Setting the Priority Based on Interface State

The tracking feature is disabled by default. When you specify the tracking option, the priority of the virtual router is changed based on the state of another interface in the switch. When the tracked interface is down, the priority of the virtual router is changed to a lower priority value. When the tracked interface is up, the priority of the virtual router is restored to its original value. You can track one of two interfaces on a switch in the Cisco MDS 9000 Family: a specified VSAN interface or a management interface.

To track the interface priority for a virtual router, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface vsan 1 switch(config-if)#	Configures a VSAN interface (VSAN 1).

	Command	Purpose
Step 3	switch(config-if)# vrrp 250 switch(config-if-vrrp)#	Creates a virtual router.
Step 4	switch(config-if-vrrp)# track interface mgmt 0 priority 2	Specifies the priority of the virtual router to be modified based on the state of the management interface.
	switch(config-if-vrrp)# no track	Disables the tracking feature.

Displaying VRRP Information

Use the **show vrrp vr** command to display configured VRRP information (see Examples 20-7 to 20-10).

Example 20-7 Displays VRRP Configured Information

```
switch# show vrrp vr 7 interface vsan 2 configuration
vr id 7 configuration
admin state down
priority 100
no authentication
advertisement-Interval 1
preempt yes
tracking interface vsan1 priority 2
protocol IP
```

Example 20-8 Displays VRRP Status Information

```
switch# show vrrp vr 7 interface vsan 2 status
vr id 7 status
MAC address 00:00:5e:00:01:07
Operational state: init
```

Example 20-9 Displays VRRP Statistics

```
switch# show vrrp vr 7 interface vsan 2 statistics
vr id 7 statistics
Become master 0
Advertisement 0
Advertisement Interval Error 0
Authentication Failure 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Authentication Type 0
Mismatch Authentication 0
Invalid Packet Length 0
```

Example 20-10 Displays VRRP Cumulative Statistics

```
switch# show vrrp statistics
Invalid checksum 0
Invalid version 0
Invalid VR ID 0
```

Clearing VRRP Statistics

Use the **clear vrrp** command to clear all the software counters for the specified virtual router (see [Example 20-11](#)).

Example 20-11 Clears VRRP Information

```
switch# clear vrrp 7 interface vsan2
switch#
```

Configuring DNS Server

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

To configure a DNS server, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ip domain-lookup	Enables the IP Domain Naming System (DNS)-based host name-to-address translation.
	switch(config)# no ip domain-lookup	Disables (default) the IP DNS-based host name-address translation and reverts to the factory default.
Step 3	switch(config)# no ip domain-name cisco.com	Disables the domain name and reverts to the factory default.
	switch(config)# ip domain-name cisco.com	Enables (default) the default domain name feature used to complete unqualified host names. Any IP host name that does not contain a domain name (that is, any name without a dot), will have the dot and cisco.com appended to it before being added to the host table.
Step 4	switch(config)# ip domain-list harvard.edu switch(config)# ip domain-list stanford.edu switch(config)# ip domain-list yale.edu	Defines a list of default domain names to complete unqualified host names, use the ip domain-list global configuration command. You can define up to 10 domain names in this list. To delete a name from a list, use the no form of this command.
	switch(config)# no ip domain-list	Deletes the defined list and reverts to factory default. No domains are configured by default.
Note	If you have not configured a domain list, the domain name that you specified with the ip domain-name global configuration command is used. If you did configure a domain list, the default domain name is not used. The ip domain-list command is similar to the ip domain-name command, except that with the ip domain-list command you can define a list of domains, each to be tried in turn.	

	Command	Purpose
Step 5	switch(config)# ip name-server 15.1.0.1 15.2.0.0	Specifies the first address (15.1.0.1) as the primary server and the second address (15.2.0.0) as the secondary sever. You can configure a maximum of six servers.
	switch(config)# no ip name-server	Deletes the configured server(s) and reverts to factory default. No server is configured by default.
	Note Alternatively, you can configure the DNS entry using the switch names (instead of IP addresses). The configured switch name automatically looks up the corresponding IP address.	

The DNS server may be dropped after two attempts due to the following reasons:

- if the IP address or the switch name is wrongly configured
- if the DNS server is not reachable due to external reasons (reasons beyond our control)


Note

When accessing a telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see [Example 20-12](#)).

Example 20-12 Displays Configured Host Details

```
switch# show hosts
Default domain is cisco.com
Domain list: ucsc.edu harvard.edu yale.edu stanford.edu
Name/address lookup uses domain service
Name servers are 15.1.0.1 15.2.0.0
```

Default Settings

[Table 20-3](#) lists the default settings for IP features.

Table 20-3 Default IPFC Settings

Parameters	Default
VSAN IP interface configuration	No IP address is assigned by default.
IP routing	Disabled.
Domain lookup	Disabled.
Domain name	Enabled.
Domain list	No domains are configured.
Name server	No servers are configured.
Virtual router	Disabled (shutdown).
Virtual router priority for switches with secondary IP address	100.

Table 20-3 *Default IPFC Settings (continued)*

Parameters	Default
Virtual router priority for switches with primary IP address	255.
Time interval between advertisement frames	1 second.
Preempting master VR	Enabled.
VRRP security authentication	No authentication.
Interface state tracking	Disabled.



Configuring FICON

Fibre Connection (FICON) interface capabilities enhances the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. Inclusion of Control Unit Port (CUP) support further enhances the MDS offering by allowing inband management of the switch from FICON processors.

This chapter includes the following sections:

- [About FICON, page 21-2](#)
- [MDS-Specific FICON Advantages, page 21-2](#)
- [FICON Port Numbering, page 21-7](#)
- [MDS FICON Prerequisites, page 21-10](#)
- [Enabling FICON, page 21-11](#)
- [Assigning Port Numbers to PortChannels, page 21-15](#)
- [Assigning FCIP Interfaces, page 21-16](#)
- [Configuring Code Page, page 21-16](#)
- [Configuring the FC ID Last Byte, page 21-16](#)
- [Configuring FICON Host Control, page 21-17](#)
- [Configuring FICON SNMP Control, page 21-18](#)
- [Automatically Saving the Running Configuration, page 21-19](#)
- [Configuring FICON Port Addresses, page 21-20](#)
- [FICON Configuration Files, page 21-22](#)
- [Port Swapping, page 21-25](#)
- [Clearing FICON Device Allegiance, page 21-26](#)
- [CUP Inband Management, page 21-26](#)
- [Displaying FICON Information, page 21-27](#)
- [Configuring Fabric Binding, page 21-33](#)
- [Displaying RLIR Information, page 21-41](#)



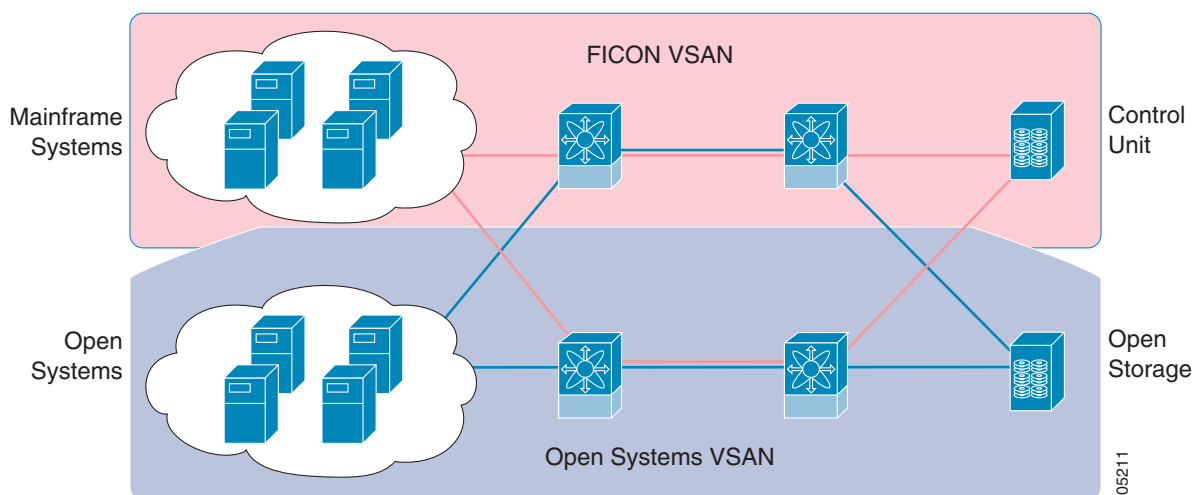
Note

FICON features can be implemented in any switch in the Cisco MDS 9000 Family running SAN-OS Release 1.3(x) or above. No hardware changes are required to configure FICON parameters.

About FICON

The Cisco MDS 9000 Family supports Fibre Channel protocol (FCP), FICON, iSCSI, and FCIP capabilities within a single, high availability platform. This solution simplifies purchasing, reduces deployment and management costs, and reduces the complex evolution to shared mainframe and open systems storage networks (see [Figure 21-1](#)).

Figure 21-1 Shared System Storage Network



FCP and FICON are different FC4 protocols and their traffic are independent of each other. If required, devices using this protocol can be isolated using VSANs.

MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches:

- [Fabric-Optimization with VSANs, page 21-3](#)
- [FCIP Support, page 21-4](#)
- [PortChannel Support, page 21-4](#)
- [VSANs for FICON and FCP Intermixing, page 21-4](#)
- [MDS-Supported FICON Features, page 21-5](#)

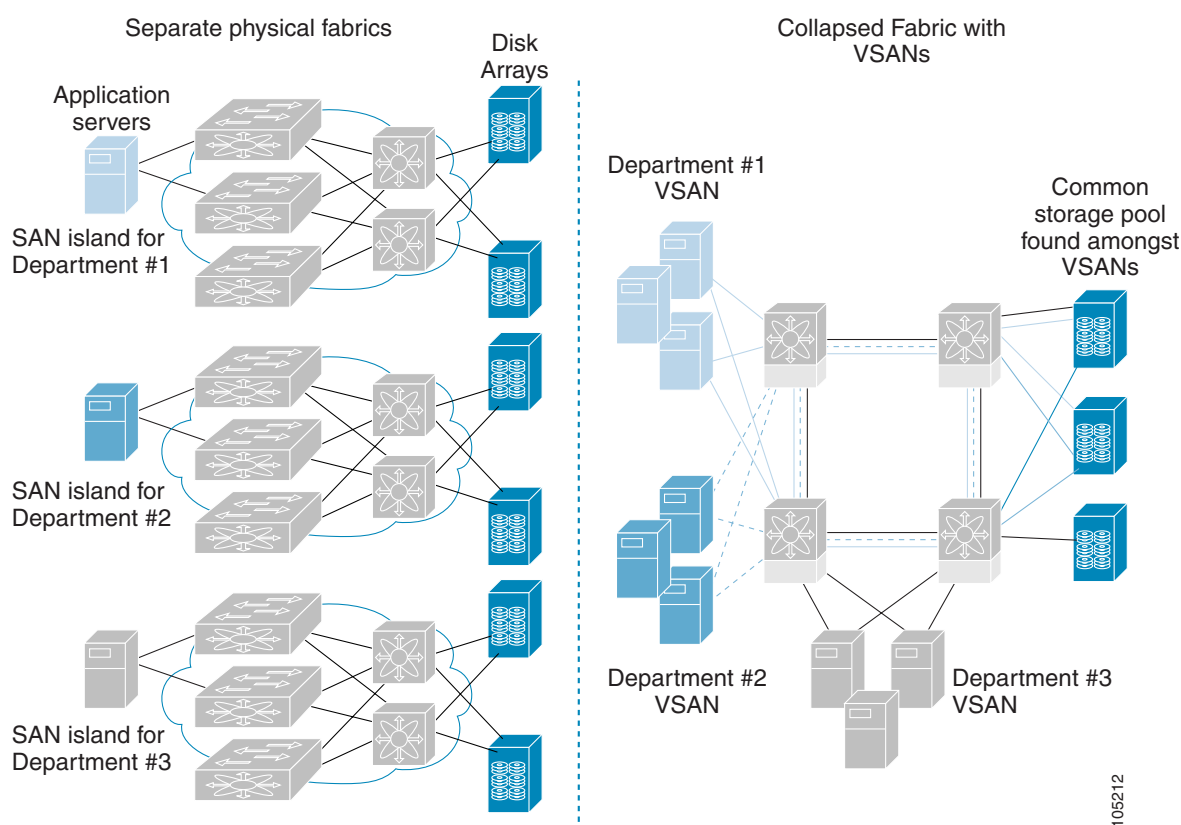
Fabric-Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. Further, the ports in each island may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can introduce greater efficiency between these physical fabric by lowering the cost of over-provisioning and reducing the number of switches to be managed.

VSANs also help you to move unused ports nondisruptively and provides a common redundant physical infrastructure (see [Figure 21-1](#)).

Figure 21-2 VSAN-specific Fabric Optimization



VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.



Note

You can configure up to 256 VSANs in any Cisco MDS switch, and you can only configure FICON in eight of these VSANs.

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and MDS 9216 switches transparently integrate Fibre Channel, FICON and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the MDS 9000 platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure and simplifying business continuance strategies.

The Cisco MDS implementation of FICON provides support for IP tunneling to efficiently consolidate SANs over WAN distances. IP tunnels enable a globally accessible storage infrastructure. Refer to the [Chapter 22, “Configuring IP Storage”](#) for further information on FCIP.

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of inter-switch links necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches. Refer to the [Chapter 12, “Configuring PortChannels”](#) for further information on PortChannels.

VSANs for FICON and FCP Intermixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex intermix environments. Multiple logical FICON, Z-Series Linux/FCP and Open-Systems FCP fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol specific fabric services, eliminating the complexity and potential instability of zone-based intermix schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Intermixed environments are addressed by the SAN-OS software. The challenge of mixing Fibre Channel Protocol (FCP) and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and Directors in the Cisco MDS 9000 Family support FCP and FICON protocol intermixing at the port level. If these protocols are intermixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

**Tip**

When creating an intermix environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

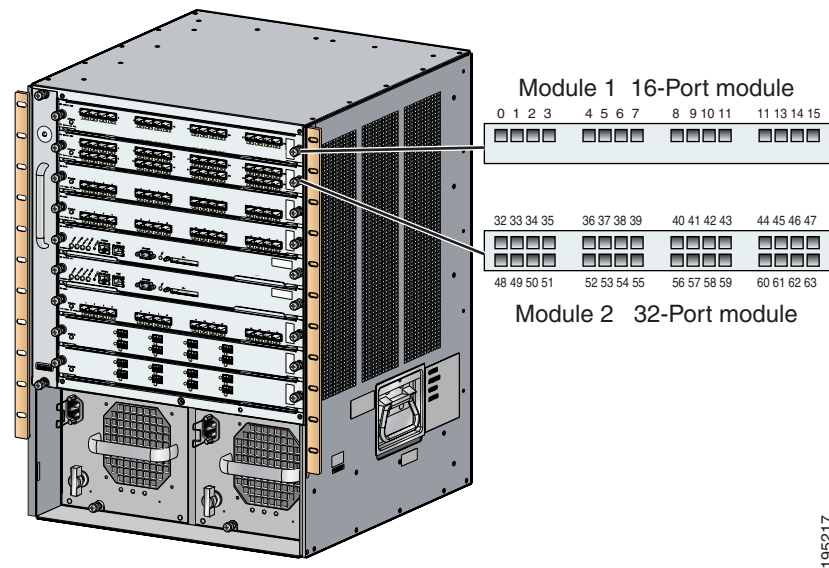
- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across all Cisco MDS 9500 Series as well as the Cisco MDS 9216 Switch. (refer to the *Cisco MDS 9500Series* and the *Cisco MDS 9216 Switch Hardware Installation Guides*).
- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 224 2/1-Gbps, autosensing FICON or Fibre Channel FCP ports in any combination in a single chassis and up to 768 Fibre Channel ports in a single rack—1.44 Tbps of internal system bandwidth ensures smooth integration of future 10-Gbps modules (see the [Chapter 5, “Configuring High Availability”](#)).
- Infrastructure protection—Common software releases infrastructure protection is available across all Cisco MDS 9000 platforms (see [Chapter 6, “Software Images”](#)).
- VSAN technology—The Cisco MDS 9000 Family introduces VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON intermix support (see [Chapter 9, “Configuring and Managing VSANs”](#)).
- Port-level configurations:
 - BB_credits for each port (see the [“Configuring Buffer-to-Buffer Credits”](#) section on [page 10-11](#)).
 - Port security for each port (see [Chapter 18, “Configuring Port Security”](#)).
 - Enable beaconing for ports and the director unit (see the [“Identifying the Beacon LEDs”](#) section on [page 10-13](#)).
- Configure an alias name, instead of the WWN, for switches and attached node devices (see [Chapter 13, “Configuring and Managing Zones”](#)).
- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN zoning, read-only zones, and VSAN-based access control (see [Chapter 16, “Configuring Switch Security”](#) and [Chapter 17, “Configuring Fabric Security”](#)).
- View the local accounting log to locate FICON events (see the [“Local AAA”](#) section on [page 16-15](#)).
- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console (see the [“CUP Inband Management”](#) section on [page 21-26](#)).
- Port address-based configurations—port name, blocked or unblocked state, and the prohibit connectivity attributes (see the [“Configuring FICON Port Addresses”](#) section on [page 21-20](#)).
- Display the following information (see [“Displaying FICON Information”](#) section on [page 21-27](#)):
 - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
 - Nodes attached to ports.
 - Port performance and statistics.
- Store and apply configuration files (see the [“FICON Configuration Files”](#) section on [page 21-22](#)).

- FICON and Open Systems Management Server features if installed (see the [“VSANs for FICON and FCP Intermixing”](#) section on page 21-4).
- Enhanced Cascading Support (see the [“CUP Inband Management”](#) section on page 21-26)
- Set the date and time on the switch (see the [“Configuring FICON Host Control”](#) section on page 21-17).
- Configure SNMP trap recipients and community names (see the [“Configuring FICON SNMP Control”](#) section on page 21-18).
- Call Home configurations—director name, location, description, and contact person (see [Chapter 23, “Configuring Call Home”](#)).
- Configure preferred domain ID, FC ID persistence, and principle switch priority (see [Chapter 24, “Configuring Domain Parameters”](#)).
- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol, decoding, and network analysis tools as well as integrated call-home capability for added reliability, faster problem resolution, and reduced service costs (see [Chapter 28, “Monitoring Network Traffic Using SPAN”](#)).
- Configure R_A_TOV, E_D_TOV (see the [“Configuring FC Timers”](#) section on page 29-2)
- Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis (see [Chapter 31, “Monitoring System Processes and Logs”](#)).
- Display and clear port-level incident alerts ([“Clearing RLIR Information”](#) section on page 21-44)

FICON Port Numbering

With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the *port number*. Port numbers are assigned based on the module and the slot in the chassis. Port numbers cannot be changed and the first port in a switch always starts with a 0 (see Figure 21-3).

Figure 21-3 Port Number in the Cisco MDS 9000 Family



The FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Even if the module is a 16-port module, 32-port numbers are assigned to that module—regardless of the module type (16-port or 32-port), the module’s physical presence in the chassis, or the port status (up or down).



Note

Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

Table 21-1 lists the port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 21-1 FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Implemented Port Numbers Allocation		Unimplemented Port Number s	Notes
		To Ports	To PortChannels		
Cisco MDS 9120 Switch	Not applicable	Ports 0 through 39	40 through 55	Ports 56 through 253 and Port 255	Only 20 ports are used.
Cisco MDS 9140 Switch	Not applicable	Ports 0 through 39	40 through 55	Ports 56 through 253 and Port 255	All 40 ports are used.

Table 21-1 FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Implemented Port Numbers Allocation		Unimplemented Port Number s	Notes
		To Ports	To PortChannels		
Cisco MDS 9216 Switch	Slot 1	Ports 0 through 31	64 through 79	Ports 80 through 253 and Port 255	Similar to a switching module
	Slot 2	Ports 32 through 63			The first 16 port numbers in a 16-port module are used and the rest remain unused.
Cisco MDS 9506 Director	Slot 1	Ports 0 through 31	Ports 128 through 143	Ports 144 through 253 and Port 255	
	Slot 2	Ports 32 through 63			
	Slot 3	Ports 64 through 95			
	Slot 4	Ports 96 through 127			
	Slot 5	None			Supervisor module are not allocated port numbers.
	Slot 6	None			
Cisco MDS 9509 Director	Slot 1	Ports 0 through 31	Ports 224 through 239	Ports 240 through 253 and Port 255	The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 2	Ports 32 through 63			
	Slot 3	Ports 64 through 95			
	Slot 4	Ports 96 through 127			
	Slot 5	None			Supervisor module are not allocated port numbers.
	Slot 6	None			
	Slot 7	Ports 128 through 159			The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 8	Ports 160 through 191			
	Slot 9	Ports 192 through 223			

Implemented and Unimplemented Ports

An implemented port refers to any port number that is available in the chassis. These numbers are identified in the [Implemented Port Numbers Allocation](#) column in [Table 21-1](#).

An unimplemented port refers to any port number that is not available in the chassis. These numbers are identified in the [Unimplemented Port Number s](#) column in [Table 21-1](#).



Tip

An unimplemented port is prohibited from communicating with an implemented port in a FICON setup and cannot be configured.

Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed if any of the following conditions apply:

- The module is not present.
- The small form-factor pluggable (SFP) port is not present.
- The port is not in a FICON-enabled VSAN.

- The port is part of a PortChannel number allocation.

For example:

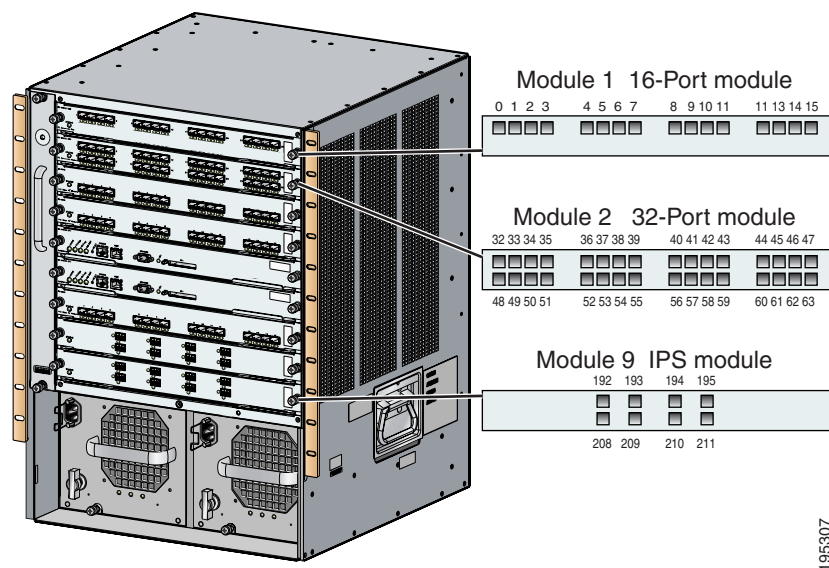
- If module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, port numbers 0 to 31 are considered uninstalled.
- If a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, port numbers 38 to 63 are considered uninstalled.
- If port number 4 (of a 16-port module in slot 1) is configured in FICON-enabled VSAN 2, then only port 4 is installed and ports 0 to 3 and 5 to 15 are uninstalled—even if they are implemented.
- If interface fc1/1 (port address = 0) is a TE port and is configured in VSANs 1 through 20—but only VSANs 2 and 3 are FICON-enabled, then port address 0 is only installed in VSAN 2 and VSAN 3.

FCIP Port Number

You must explicitly configure FCIP port numbers. The port address for FCIP ports are configured to the range of numbers that you can use are restricted to the port numbers available in the IPS modules slot. If an IPS module is in Slot 9 in a Cisco MDS 9509 Director, the available range of port numbers is 192 through 223. The FCIP interface can be assigned any port number that is available within that range.

For example, if the FCIP port is bound to GigabitEthernet interface 9/1, the assigned FCIP port numbers can be 192, 193, 194, 195, 208, 209, 210, or 211 (see [Figure 21-3](#)).

Figure 21-4 FCIP Port Numbers in the Cisco MDS 9000 Family



Note

Gigabit Ethernet ports do not have a corresponding mapping to the FICON port number concept.

Use the **show fcip portnumber** command to view the list of available port numbers for a specified module.

Port Numbering Summary

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers are VSAN independent—Fibre Channel port numbers do not change based on VSANs or TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- Physical ports in a PortChannel become uninstalled. These port numbers are not applied in the FICON configurations, but the PortChannel configuration is applied to the physical ports.
- A FCIP tunnel must be explicitly associated with a FICON port number—If the port number is not assigned for PortChannels or for FCIP tunnels, the associated ports will not come up (see the [“FCIP Port Number” section on page 21-9](#)).
- iSCSI and virtualization ports are not exposed to FICON and do not have FICON port number associations.

Port Addresses

By default, port numbers are the same as port addresses. The default port address changes when you issue the port swap command at any time. When you issue this command, the port addresses are swapped (see the [“Port Swapping” section on page 21-25](#)).

FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 and can be configured (see the [“Configuring the FC ID Last Byte” section on page 21-16](#)).

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are flapped to switch from the dynamic to static FC IDs and vice versa.

**Note**

You cannot configure persistent FC IDs in FICON-enabled VSANs.

MDS FICON Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature (see the [“The Default Zone” section on page 13-9](#)).
- Enable in-order delivery on the switch (see the [“In-Order Delivery” section on page 19-10](#)).
- Enable (and if required, configure) fabric binding on the VSAN (see the [“Configuring Fabric Binding” section on page 21-33](#)).
- Verify that conflicting persistent FC IDs do not exist in the switch (see [Chapter 24, “Configuring Domain Parameters”](#)).

- Verify that the configured domain ID and requested domain ID match (see [Chapter 24, “Configuring Domain Parameters”](#)).
- Add the CUP (area FE) to the zone, if you are using zoning (see the [“CUP Inband Management” section on page 21-26](#)).

If any of these requirements are not met, the FICON feature cannot be enabled.

Enabling FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on a per VSAN basis in one of three ways:

- By using the automated **setup ficon** command (see the [“Setting Up a Basic FICON Configuration” section on page 21-11](#)).
- Manually addressing each prerequisite (see the [“Manually Enabling FICON” section on page 21-14](#)).
- By using the Device Manager (refer to the *Cisco MDS 9000 Family Fabric Manager User Guide* for further information).

Effects of Enabling FICON

When you enable the FICON feature in Cisco MDS switches, the following apply:

- The IPL configuration file is automatically created (see the [“FICON Configuration Files” section on page 21-22](#)).
- You cannot disable in-order delivery, fabric binding, or static domain ID configurations.

Setting Up a Basic FICON Configuration

This section steps you through the procedure to setup FICON on a specified VSAN in a Cisco MDS 9000 Family switch.



Note

Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point.



Tip

If you do not wish to answer a previously-configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example switch name), the switch uses what is previously configured and skips to the next question.

To enable and setup FICON, follow these steps.

Step 1 Issue the **setup ficon** command at the EXEC command mode.

```
switch# setup ficon
      --- Ficon Configuration Dialog ---
```

This setup utility will guide you through basic Ficon Configuration

on the system.

Press Enter if you want to skip any dialog. Use ctrl-c at anytime to skip all remaining dialogs.

Step 2 Enter **yes** (the default is **yes**) to enter the basic FICON configuration setup.

Would you like to enter the basic configuration dialog (yes/no) [yes]: **yes**

The FICON setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt, to end the configuration process.

Step 3 Enter the VSAN number for which FICON should be enabled:

Enter vsan [1-4093]:2

Step 4 Enter **yes** (the default is **yes**) to confirm your VSAN choice:

Enable ficon on this vsan? (yes/no) [yes]: **yes**



Note At this point, the software creates the new VSAN if it does not already exist.

Step 5 Enter the domain ID number for the specified VSAN:

Configure domain-id for this ficon vsan (1-239):2

Step 6 Enter **yes** (the default is **no**) to set up FICON in cascaded mode. If you enter **no** at this point, skip to the Step 7 (see [“CUP Inband Management” section on page 21-26](#) for further details).

Would you like to configure ficon in cascaded mode: (yes/no) [no]: **yes**

a. Assign the peer WWN for the FICON: CUP.

Configure peer wwn (hh:hh:hh:hh:hh:hh:hh:hh): **11:00:02:01:aa:bb:cc:00**

b. Assign the peer domain ID for the FICON: CUP

Configure peer domain (1-239) :4

c. Enter **yes** if you wish to configure additional peers (and repeat Steps 6a and 6b). Enter **no**, if you do wish to configure additional peers.

Would you like to configure additional peers: (yes/no) [no]: **no**

Step 7 Enter **no** (the default is **no**) to deny SNMP permission to modify existing port connectivity parameters (see [“Configuring FICON SNMP Control” section on page 21-18](#) for further details).

Enable SNMP to modify port connectivity parameters? (yes/no) [no]: **no**

Step 8 Enter **no** (the default is **no**) to disable the host (mainframe) to modify the port connectivity parameters, if required (see [“Configuring FICON Host Control” section on page 21-17](#) for further details).

Disable Host from modifying port connectivity parameters? (yes/no) [no]: **no**

Step 9 Enter **yes** (the default is **yes**) to enable the **active equals saved** feature (see [“Automatically Saving the Running Configuration” section on page 21-19](#) for further details).

Enable active=saved? (yes/no) [yes]: **yes**

Step 10 Enter **no** (the default is **yes**) if you do not wish to configure additional FICON VSANs.

Would you like to configure additional ficon vsans (yes/no) [yes]: **no**

Step 11 Review and edit the configuration that you have just entered.

Step 12 Enter **no** (the default is **no**) if you are satisfied with the configuration.



Note For documentation purposes, the following configuration displays three VSANs with different FICON settings. These settings provide a sample output for different FICON scenarios.

The following configuration will be applied:

```
fcdomain domain 2 static vsan 1
fcdomain restart disruptive vsan 1
fabric-binding database vsan 1
swwn 11:00:02:01:aa:bb:cc:00 domain 4
fabric-binding activate vsan 1
zone default-zone permit vsan 1
ficon vsan 1
no host port control
```

```
fcdomain domain 3 static vsan 2
fcdomain restart disruptive vsan 2
fabric-binding activate vsan 2 force
zone default-zone permit vsan 2
ficon vsan 2
no host port control
no active equals saved
```

```
vsan database
vsan 3
fcdomain domain 5 static vsan 3
fcdomain restart disruptive vsan 3
fabric-binding activate vsan 3 force
zone default-zone permit vsan 3
ficon vsan 3
no snmp port control
no active equals saved
```

Would you like to edit the configuration? (yes/no) [no]: **no**

Step 13 Enter **yes** (the default is **yes**) to use and save this configuration. The implemented commands are displayed. After FICON is enabled for the specified VSAN, you are returned to the EXEC mode switch prompt.

Use this configuration and save it? (yes/no) [yes]: **yes**

```
`fcdomain domain 2 static vsan 1`
`fcdomain restart disruptive vsan 1`
`fabric-binding database vsan 1`
`swwn 11:00:02:01:aa:bb:cc:00 domain 4`
Error: WWN is invalid
*** submode cmd exec error ***
`fabric-binding activate vsan 1`
`zone default-zone permit vsan 1`
`ficon vsan 1`
`no host port control`

`fcdomain domain 3 static vsan 2`
`fcdomain restart disruptive vsan 2`
`fabric-binding activate vsan 2 force`
`zone default-zone permit vsan 2`
`ficon vsan 2`
`no host port control`
`no active equals saved`
```

```

`vsan database`
`vsan 3`
`fcdomain domain 5 static vsan 3`
`fcdomain restart disruptive vsan 3`
`fabric-binding activate vsan 3 force`
`zone default-zone permit vsan 3`
`ficon vsan 3`
`no snmp port control`
`no active equals saved`

switch#

```



Note If a new VSAN is created, two additional commands are displayed— **vsan database** and **vsan number**.

Manually Enabling FICON

To manually enable FICON on a VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# vsan database switch(config-vsan-db)# vsan 2 switch(config-vsan-db)# do show vsan usage 2 vsan configured configured vsans:1-2 vsans available for configuration:3-4093 switch(config-vsan-db)# exit	Enables VSAN 2.
Step 3	switch(config)# in-order-guarantee	Activates in-order delivery for the entire switch. See the Chapter 19, “Configuring Fibre Channel Routing Services and Protocols.”
Step 4	switch(config)# fcdomain domain 2 static vsan 2	Configures the domain ID for VSAN 2. See the Chapter 24, “Configuring Domain Parameters.”
Step 5	switch(config)# fcdomain restart disruptive vsan 2 Note This is an optional step—apply this step only if the current running domain ID and the static domain ID are not the same. Use the show fcdomain vsan command to verify these IDs	Restarts the VSAN to ensure that the configured domain ID is the same as the runtime domain ID. See the Chapter 24, “Configuring Domain Parameters.”

	Command	Purpose
Step 6	<pre>switch(config)# do show fcdomain vsan 2 The local switch is the Principal Switch. Local switch run time information: State: Stable Local switch WWN: 20:02:00:05:30:00:2a:1f Running fabric name: 20:02:00:05:30:00:2a:1f Running priority: 128 Current domain ID: 0x02(2) Local switch configuration information: State: Enabled ... Configured domain ID: 0x02(2) (static) Principal switch run time information: ...</pre>	<p>Verifies that the configured domain ID and requested domain ID match.</p> <p>See the Chapter 24, “Configuring Domain Parameters.”</p>
Step 7	<pre>switch(config)# fabric-binding activate vsan 2 force</pre>	<p>Activates fabric binding on VSAN 2.</p> <p>See the “Configuring Fabric Binding” section on page 21-33.</p>
Step 8	<pre>switch(config)# zone default-zone permit vsan 2</pre>	<p>Sets the default zone to permit for VSAN 2.</p> <p>See the “CUP Inband Management” section on page 21-26.</p>
Step 9	<pre>switch(config)# ficon vsan 2</pre>	Enables FICON on VSAN 2.
	<pre>switch(config)# no ficon vsan 6</pre>	Disables the FICON feature on VSAN 6.

Assigning Port Numbers to PortChannels

You can associate a PortChannel with a FICON port number to bring up that port.

To associate a PortChannel with a FICON port number, follow these steps:

	Command	Purpose
Step 1	<pre>switch# config t switch(config)#</pre>	Enters configuration mode.
Step 2	<pre>switch(config)# interface Port-channel 1 switch(config-if)#</pre>	Enters the PortChannel interface configuration mode.
Step 3	<pre>switch(config-if)# ficon portnumber 234</pre>	Assigns the FICON port number to the selected PortChannel port.

Assigning FCIP Interfaces

You can associate a FCIP interface with a FICON port number to bring up that interface.

To associate a FCIP interface with a FICON port number, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch1(config)# interface fcip 51 switch1(config-if)#	Creates a FCIP interface (51).
Step 3	switch(config-if)# ficon portnumber 208	Assigns the FICON port number to the selected FCIP interface.

Configuring Code Page

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Use the **code-page** command to configure the EBCDIC format. Refer to your mainframe documentation for details on the code page options. Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.



Tip

This is an optional configuration. If you are not sure of the EBCDIC format to be used, we recommend retaining the **us-canada** (default) option.

To configure the **code-page** option in a VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# code-page italy	Configures the italy EBCDIC format.
	switch(config-ficon)# no code-page	Reverts to the factory default of using the us-canada EBCDIC format.

Configuring the FC ID Last Byte

FICON requires the last byte of the fabric address to be the same for all allocated FC IDs. By default, this value is set to 0. You can change the default by issuing the **fcid-last-byte** command.

To assign the last byte for the FC ID, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.

	Command	Purpose
Step 3	<code>switch(config-ficon)# fcid-last-byte 12</code>	Assigns the last byte FC ID for the fabric address.
	<code>switch(config-ficon)# no fcid-last-byte 3</code>	Removes the configured last byte FC ID for the fabric address and reverts to the factory default of 0.

Configuring FICON Host Control

Use the **host** command to configure the switch offline state, the mainframe access control parameters, and the mainframe time stamp parameters.

Setting the Switch State

Use the **host control switch offline** command to allow the host to move the switch to an offline state. When you issue this command, all ports are shut down.

To move the switch to an offline state, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# ficon vsan 2</code> <code>switch(config-ficon)#</code>	Enables FICON on VSAN 2.
Step 3	<code>switch(config-ficon)# no host control switch offline</code>	Prohibits mainframe users from moving the switch to an offline state.
	<code>switch(config-ficon)# host control switch offline</code>	Allows the host to move the switch to an offline state and shut down the ports (default).

Controlling Mainframe Access

Use the host port control command to permit mainframe users to configure FICON parameters. By default, mainframe users are allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

To configure mainframe access, follow these steps:

	Command	Purpose
Step 1	<code>switch# config t</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# ficon vsan 2</code> <code>switch(config-ficon)#</code>	Enables FICON on VSAN 2.
Step 3	<code>switch(config-ficon)# no host port control</code>	Prohibits mainframe users to configure FICON parameters in the Cisco MDS switch.
	<code>switch(config-ficon)# host port control</code>	Allows mainframe users to configure FICON parameters in the Cisco MDS switch (default).

Setting the Time Stamp

Each VSAN in a Cisco MDS switch represents a virtual director. The clock and time present in each virtual director can be different. To maintain separate clocks for each VSAN, the SAN-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. Whenever the host (mainframe) sets the time, the SAN-OS software updates this difference between the clocks. When a mainframe reads the clock, it computes the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe. The VSAN-clock's current time is reported to the user in the output of **show ficon vsan** *vsan-id*, **show ficon**, and **show accounting log** commands.

By default, the clock in each VSAN is the same as the switch hardware clock. Mainframe users are allowed to change the VSAN-clock.

To configure host control, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# no host set-timestamp	Prohibits mainframe users from changing the VSAN-specific clock.
	switch(config-ficon)# host set-timestamp	Allows the host to set the clock on this switch (default).

Clearing Time Stamps

Use the **clear ficon vsan** *vsan-id* **timestamp** command in EXEC mode to clear the VSAN-clock for the specified VSAN.

```
switch# clear ficon vsan 20 timestamp
```

Configuring FICON SNMP Control

By default, SNMP users can configure FICON parameters through the Fabric Manager application. You can prohibit this access, if required, by issuing the **no snmp port control** command.

To configure SNMP control, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# no snmp port control	Prohibits SNMP users from configuring FICON parameters.
Step 4	switch(config-ficon)# snmp port control	Allows SNMP users to configure FICON parameters (default).

Automatically Saving the Running Configuration

When **active equals saved** is enabled, any FICON changes to the block, prohibit or port address name are immediately written to the IPL file (see the “[FICON Configuration Files](#)” section on page 21-22). You do not need to issue a **copy running startup** command to save this configuration.

If **active equals saved** is enabled in at least one FICON-enabled VSAN in the fabric, the following apply:

- All non-FICON configuration changes are automatically saved to persistent storage when a configuration is changed.
- All FICON configuration changes (block, prohibit, port address name, and host control) are not saved to persistent storage. They are only saved in a FICON-enabled VSAN—if the **active equals saved** feature is enabled in that VSAN.



Note

If **active equals saved** is enabled, the SAN-OS software ensures that you do not have to perform the **copy running startup** command for the FICON configuration as well. If your switch or fabric consists of multiple FICON-enabled VSANs, and one of these VSANs have **active equals saved** enabled, changes made to the non-FICON configuration results in all configurations being saved to the startup configuration.

To automatically save the running configuration, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# active equals saved	Enables the automatic save feature for all VSANs in the switch or fabric.
	switch(config-ficon)# no active equals saved	Disables the automatic save feature for this VSAN.

Configuring FICON Port Addresses

You can configure FICON port addresses on a per-port address basis in any switch in the Cisco MDS 9000 Family.

Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

Blocking Ports

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic is not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit Off-Line State (OLS) primitive sequence on a blocked port.



Note

The **shutdown/no shutdown** port state is independent of the **block/no block** port state. If a port is shutdown, unblocking that port will not initialize the port.



Caution

You cannot block or prohibit CUP port (0XFE).

To block or unblock port addresses in a VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# portaddress 1 - 5 switch(config-ficon-portaddr)#	Selects port address 1 to 5 for further configuration.
Step 4	switch(config-ficon-portaddr)# block	Disables a range of port addresses and retains it in the operationally down state.
	switch(config-ficon-portaddr)# no block	Enables the selected port address and reverts to the factory default of the port address not being blocked.

Prohibiting Ports

To prevent implemented ports (see the [“Implemented and Unimplemented Ports”](#) section on page 21-8) from talking to each other, you can configure two or more ports to be in a prohibited state. If you prohibit ports, the specified ports are prevented from communicating with each other.



Note

Unimplemented ports are always prohibited.

Prohibit configurations are always symmetrically applied—if you prohibit Port 0 from talking to Port 15, Port 15 is automatically prohibited from talking to Port 0.

Physical Fibre Channel port blocks continue to transmit OLS and NOS protocols on that link.

To prohibit port addresses in a VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# portaddress 7 switch(config-ficon-portaddr)#	Selects port address 7 for further configuration.
Step 4	switch(config-ficon-portaddr)# prohibit portaddress 3-5 switch(config-ficon-portaddr)# no prohibit portaddress 5	Prohibits port address 7 in VSAN 2 from talking to ports 3, 4, and 5. Removes port address 5 from a previously-prohibited state.

Assigning Port Address Names

To assign a port address name, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# portaddress 7 switch(config-ficon-portaddr)#	Selects port address 7 for further configuration.
Step 4	switch(config-ficon-portaddr)# name SampleName switch(config-ficon-portaddr)# no name SampleName	Assigns a name to the port address. Deletes a previously configured port address name.

FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBMTM. These files can be read and written by IBM hosts using the inband CUP protocol. Additionally, you can use the Cisco MDS CLI or FM applications to operate these FICON configuration files

**Note**

Multiple FICON configuration files with the same name can exist in the same switch, provide they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always uses the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled on a VSAN.

FICON configuration files contain the following configuration for each implemented port address:

- Host control
- Block
- Prohibit mask
- Port address name

**Note**

Refer to the [“Working with Configuration Files” section on page 4-23](#) for details on the normal configuration files used by Cisco MDS switches. This configuration file includes FICON enabled attribute for a VSAN, port number mapping for port channels and FCIP interfaces, port number to port address mapping ([“Port Swapping” section on page 21-25](#)), port and trunk allowed VSAN configuration for ports, in-order guarantee, configuring static domain ID, and fabric binding configuration ([“Configuring Fabric Binding” section on page 21-33](#)).

Writing to the IPL file

When configuring FICON you will predominantly be working with FICON-related configurations. The non-FICON configurations are used to initially configure the switch. You can save FICON configuration files using two methods: snapshot of running configuration or automatically saving the running configuration.

**Caution**

When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

Snapshot of Running Configuration

When you issue the **copy running startup** command, MDS saves a snapshot of the complete running configuration to the startup configuration. This includes saving the FICON running configuration to the IPL file for each VSAN on which FICON is enabled.

Accessing FICON Configuration Files

Only one user can access the configuration file at any given time:

- While this file is being accessed by user 1, user 2 cannot access this file.
- When user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 has been inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host SNMP CLI user who is permitted to access the switch. The locking mechanism in the SAN-OS software restricts access to one user at a time per file. This lock applies to newly-created files and previously-saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

Applying the FICON Configuration Files

The configuration from the saved files can be applied to the running configuration by using the **ficon vsan number apply file filename** command. For example:

```
switch# ficon vsan 2 apply file SampleFile
```

This command reads the file content and applies the changes to the running configuration. However if **active equals saved** is enabled, the SAN-OS software also saves the changes to the IPL file.

When an MDS switch is booting up with saved configuration, if FICON is enabled on a VSAN, the IPL configuration file is applied automatically by the SAN-OS software after the switch initialization is completed.

Editing FICON Configuration Files

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to 8 alphanumeric characters.

To edit the contents of a specified FICON configuration file, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# ficon vsan 2 switch(config-ficon)#	Enables FICON on VSAN 2.
Step 3	switch(config-ficon)# file IplFile1 switch(config-ficon-file)#	Accesses the FICON configuration file called IplFile1 for VSAN 2. If this file does not exist, it is created.
		Note All FICON file names are restricted to 8 alphanumeric characters.
	switch(config-ficon)# no file IplFileA	Deletes a previously-created FICON configuration file.

	Command	Purpose
Step 4	switch(config-ficon-file)# portaddress 3 switch(config-ficon-file-portaddr)#	Enters the submode for port address 3 to edit the contents of the configuration file named IplFile1. Note The running configuration is not applied to the current configuration. The configuration is only applied when the ficon vsan number apply file filename command is issued.
Step 5	switch(config-ficon-file-portaddr)# prohibit portaddress 5	Edits the content of the configuration file named IplFile1 by prohibiting port address 5 from accessing port address 3.
Step 6	switch(config-ficon-file-portaddr)# block	Edits the content of the configuration file named IplFile1 by blocking a range of port addresses and retaining them in the operationally down state.
Step 7	switch(config-ficon-file-portaddr)# name P3	Edits the content of the configuration file named IplFile1 by assigning the name P3 to port address 3. If the name did not exist, it is created. If it existed, it is overwritten.

Copying FICON Configuration Files

Use the **ficon vsan vsan-id copy file exiting-file-name save-as-file-name** command in EXEC mode to copy an existing FICON configuration file.

```
switch# ficon vsan 20 copy file IPL IPL3
```

You can see the list of existing configuration files by issuing the **show ficon vsan vsan-id** command.

```
switch# show ficon vsan 20
Ficon information for VSAN 20
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Disabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Enabled
  Number of implemented ports are 240
  Key Counter is 5
  FCID last byte is 0
  Date/Time is same as system time (Wed Dec 3 20:10:45.924591 2003)
  Device Allegiance not locked
  Codepage is us-canada
Saved configuration files
  IPL
  IPL3
```


Port Swapping

The port swap FICON feature is only provided for maintenance purposes and is supported in all switches in the Cisco MDS 9000 Family support this feature.

The **ficon swap portnumber** *old-port-number new port-number* command causes all configuration associated with *old-port-number* and *new port-number* to be swapped, including VSAN configurations. This command is only associated with the two ports in concerned. You must issue this VSAN-independent command from the EXEC mode.



Tip

If **active equals saved** is enabled on any FICON VSAN then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly issue the **copy running startup** command immediate after swapping the ports.

MDS switches also allow port swapping for non-existent ports as specified below:

- Only FICON-specific configurations (port number to port address mapping) is swapped
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.

Once you issue the **ficon swap portnumber** *old-port-number new-port-number* command, the switch automatically performs the following actions:

- Shuts down both the old and new ports
- Swaps the port configuration
- Initializes the port shut down if you specified the **after swap noshut** option after the *new-port-number*. Otherwise, you must explicitly issue the **no shutdown** to resume traffic.

To swap physical Fibre Channel ports, follow these steps:

- Step 1** Issue the **ficon swap portnumber** *old-port-number new-port-number* command in EXEC mode. The specified ports are operationally shut down.
- Step 2** Physically swap the front panel port cables between the two ports.
- Step 3** Issue the **no shutdown** command on each port to enable traffic flow.



Note

If you specify the **ficon swap portnumber** *old-port-number new-port-number* **after swap noshut** command, the ports will automatically be initialized.

Port Swapping Guidelines

Be sure to follow these guidelines when using the FICON port swap feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.

- Before performing a port swap, the SAN-OS software performs a compatibility check. If the two ports have incompatible configuration, the port swap is rejected with an appropriate reason code. For example, if a port with BB_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB_credits is allowed (not a configurable parameter), the port swap operation is rejected.
- If ports have default values (for some incompatible parameters), then port swap is allowed to go through and the ports retain their default values.

**Note**

The 32-port module guidelines also apply for port swapping configurations (see the [“Configuring 32-port Switching Modules and Host-Optimized Ports”](#) section on page 10-8).

Clearing FICON Device Allegiance

FICON requires serialization of access between multiple mainframes, CLI, and SNMP sessions. You can maintain this access on Cisco MDS 9000 Family switches by controlling device allegiance for the currently-executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available. You can clear the current device allegiance by issuing the **clear ficon vsan vsan-id allegiance** command in EXEC mode.

```
switch# clear ficon vsan 1 allegiance
```

**Caution**

This command aborts the currently-executing session.

CUP Inband Management

The Control Unit Port (CUP) protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for inband management using the IBM S/A OS/390 I/O operations console.

**Note**

The CUP specification is proprietary to IBM.

CUP is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the MDS switches.

Host communication includes control functions like blocking/unblocking ports, as well as monitoring and error reporting functions.

To place the CUP in a zone, follow these steps.

Step 1 Set the default zone to permit for the required VSAN.

```
switch# config t
switch(config)# zone default-zone permit vsan 20
```

if you are not using the zoning feature (see the [“The Default Zone”](#) section on page 13-9).

- Step 2** Issue the `show fcns database` command for the required VSAN and obtain the required FICON CUP WWN.

```
switch# show fcns database vsan 20
```

```
VSAN 20:
```

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x0d0d00	N	50:06:04:88:00:1d:60:83	(EMC)	FICON:CU
0x0dfe00	N	25:00:00:0c:ce:5c:5e:c2	(Cisco)	FICON:CUP
0x200400	N	50:05:07:63:00:c2:82:d3	(IBM)	scsi-fcp FICON:CU f..
0x200800	N	50:05:07:64:01:40:15:0f	(IBM)	FICON:CH
0x20fe00	N	20:00:00:0c:30:ac:9e:82	(Cisco)	FICON:CUP

```
Total number of entries = 5
```



Note

If more than one FICON:CUP WWN exists in this fabric, be sure to add all the FICON:CUP WWN PWWNs to the required zone. The example above, displays multiple FICON:CUP occurrences to indicate a cascade configuration.

- Step 3** Add the identified FICON:CUP WWN to the zone database.

```
switch(config)# zone name Zone1 vsan 20
switch(config-zone)# member pwwn 25:00:00:0c:ce:5c:5e:c2
```

Displaying FICON Information

Use the `show` commands to display all FICON information configured on this switch (see Examples 21-1 to 21-13).

Example 21-1 Displays Configured FICON Information

```
switch# show ficon
Ficon information for VSAN 20
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Enabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Disabled
  Number of implemented ports are 240
  Key Counter is 73723
  FCID last byte is 0
  Date/Time is set by host to Sun Jun 26 00:04:06.991999 1904
  Device allegiance is locked by Host
  Codepage is us-canada
  Saved configuration files
    IPL
    _TSIRN00
```

In Example 21-1 the `clock alert mode is enabled` output confirms that you will receive an alert to indicate any change in the clock settings. Similarly the `user alert mode is enabled` output confirms that you will receive an alert to indicate any change in the FICON configuration.

Example 21-2 Displays Port Address Information

```

switch# show ficon vsan 2 portaddress
Port Address 1 is not installed in vsan 2
  Port number is 1, Interface is fc1/1
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255

Port Address 2 is not installed in vsan 2
  Port number is 2, Interface is fc1/2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255

...
Port Address 239 is not installed in vsan 2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255

Port Address 240 is not installed in vsan 2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255

```

Example 21-3 Displays Port Address Information in a Brief Format

```

switch# show ficon vsan 2 portaddress 50-55 brief
-----
Port   Port   Interface      Admin   Status      Oper   FCID
Address Number                                     Blocked  Mode
-----
50     50     fc2/18         on      fcotAbsent   --     --
51     51     fc2/19         off     fcotAbsent   --     --
52     52     fc2/20         off     fcotAbsent   --     --
53     53     fc2/21         off     fcotAbsent   --     --
54     54     fc2/22         off     notConnected --     --
55     55     fc2/23         off     up           FL     0xea0000
56     55     off            off     up           FL     0xea0000

```

In [Example 21-3](#), the interface column is populated with the corresponding interface if the port number is installed. If the port number is uninstalled, this space remains blank.

Example 21-4 Displays Port Address Counter Information

```

switch# show ficon vsan 20 portaddress 8 counters
Port Address 8(0x8) is up in vsan 20
  Port number is 8(0x8), Interface is fc1/8
  Version presented 1, Counter size 32b
  242811 frames input, 9912794 words
    484 class-2 frames, 242302 class-3 frames
    0 link control frames, 0 multicast frames
    0 disparity errors inside frames
    0 disparity errors outside frames
    0 frames too big, 0 frames too small
    0 crc errors, 0 eof errors
    0 invalid ordered sets
    0 frames discarded c3
    0 address id errors
  116620 frames output, 10609188 words

```

```

    0 frame pacing time
    0 link failures
    0 loss of sync
    0 loss of signal
    0 primitive seq prot errors
    0 invalid transmission words
    1 lrr input, 0 ols input, 5 ols output
    0 error summary

```

Example 21-5 *Displays the Contents of the Specified FICON Configuration File*

```

switch# show ficon vsan 2 file IplFile1
switch# show ficon vsan 3 file IPL
FICON configuration file IPL          in vsan 3
  Port address 1
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

  Port address 2
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

  Port address 3
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

  Port address 4
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

  ...
  Port address 80
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

  Port address 254
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

```

Example 21-6 *Displays All FICON Configuration Files*

```

switch# show ficon vsan 2
Ficon information for VSAN 2
  Ficon is enabled
  VSAN is active
  Host control is Enabled
  Host offline control is Enabled
  Clock alert mode is Disabled
  User alert mode is Disabled
  SNMP control is Disabled
  Active=Saved is Disabled
  Number of implemented ports are 240
  Key Counter is 9
  FCID last byte is 0
  Date/Time is same as system time(Sun Dec 14 01:26:30.273402 1980)
  Device Allegiance not locked

```

```
Codepage is us-canada
Saved configuration files
IPL
IPLFILE1
```

Example 21-7 Displays the Specified Port Addresses for a FICON Configuration File

```
switch# show ficon vsan 2 file SampleFile portaddress 1-3
switch# show ficon vsan 2 file iplfile1 portaddress 1-7
FICON configuration file IPLFILE1 in vsan 2
  Port address 1
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,241-253,255

  Port address 2
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,241-253,255

  Port address 3
    Port name is P3
    Port is not blocked
    Prohibited port addresses are 0,241-253,255
...
  Port address 7
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,241-253,255
```

If FICON is not enabled on a VSAN, you will not be able to view the port address information for that VSAN. 21-8 shows two command output when FICON is Enabled and Disabled.

Example 21-8 Displays the Specified Port Address When FICON is Enabled

```
switch# show ficon vsan 1 portaddress 55
FICON not enabled
switch# show ficon vsan 2 portaddress 55
Port Address 55 is not installed in vsan 2
  Port number is 55, Interface is fc2/23
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255
  Admin port mode is FL
  Port mode is FL, FCID is 0xea0000
```

If the port is blocked, the show ficon vsan number portaddress number command displays the blocked state of the port. If a specific port is prohibited, this command also displays the specifically prohibited (3) port along with the ports that are prohibited by default (0, 241 to 253, and 255). If a name is assigned that name is also displayed (see Example 21-9).

Example 21-9 Displays Two Port Addresses Configured with Different States

```
switch# show ficon vsan 2 portaddress 2
Port Address 2(0x2) is not installed in vsan 2
  Port number is 2(0x2), Interface is fc1/2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255(0,0xf1-0xfd,0xff)
  Admin port mode is auto
```

```

Peer was type model manufactured by

switch# show ficon vsan 2 portaddress 1
Port Address 2(0x2) is not installed in vsan 2
Port number is 2(0x2), Interface is fc1/2
Port name is
Port name is SampleName
Port is admin blocked
Prohibited port addresses are 0,241-253,255(0,0xf1-0xfd,0xff)
Admin port mode is auto
Peer was type model manufactured by

```

Example 21-10 Displays Control Unit Information

```

switch# show ficon control-device sb3
Control Unit Image:0x80b9c2c
VSAN:20 CU:0x20fe00 CUI:0 CUD:0 CURLP:(nil)
ASYNC LP:(nil) MODE:1 STATE:1 CQ LEN:0 MAX:0
PRIMARY LP: VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0
ALTERNATE LP: VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0

Logical Path:0x80b9fb4
VSAN:20 CH:0x200600 CHI:15 CU:0x20fe00 CUI:0 STATE:1 FLAGS:0x1
LINK: OH:0x0 OC:0x0 IH:0x0 IC:0x0
DEV: OH:0x0 OC:0x0 IH:0x0 IC:0x0
SENSE: 00 00 00 00 00 00 00 00 46
        30 20 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
IUI:0x0 DHF:0x0 CCW:0x0 TOKEN:0x0 PCCW:0x0 FCCW:0x0 PTOKEN:0x0 FTOKEN:0x0
CMD:0x0 CCW_FLAGS:0x0 CCW_COUNT:0 CMD_FLAGS:0x0 PRIO:0x0 DATA_COUNT:0
STATUS:0x0 FLAGS:0x0 PARAM:0x0 QTP:0x0 DTP:0x0
CQ LEN:0 MAX:0 DESTATUS:0x0

```

Example 21-11 Displays the History Buffer for the Specified VSAN

```

switch# show ficon vsan 20 director-history
Director History Buffer for vsan 20

```

```

-----
Key Counter          Ports Address
                    Changed
-----
74556                43
74557                44
74558                45
74559                46
74560                47
74561                48
74562                49
74563                50
74564                51
74565                52
74566                53
74567                54
74568                55
74569                56
74570                57
74571                58
74572                59
74573                60

```

74574	61
74575	62
74576	63
74577	64
74578	
74579	
74580	1-3, 5, 10, 12, 14-16, 34-40, 43-45, 47-54, 56-57, 59-64
74581	3, 5
74582	64
74583	
74584	1-3, 10, 12, 14-16, 34-40, 43-45, 47-54, 56-57, 59-64
74585	1
74586	2
74587	3

In [Example 21-11](#), the key counter (a 32-bit value) is incremented when a FICON-related configuration is changed. Host programs can increment this value at the start of the channel program and then perform operations on multiple ports. The director history buffer keeps a log of which port address configuration was changed for each key-counter value.

Example 21-12 Displays the Running Configuration Information

```
switch# show running-config
...
ficon vsan 2
portaddress 1
block
name SampleName
prohibit portaddress 3
portaddress 3
prohibit portaddress 1
file IPL
```

Example 21-13 Displays the Startup Configuration

```
switch# show startup-config
...
ficon vsan 2
file IPL
```


Configuring Fabric Binding

The SAN-OS 1.3(x) fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis and can only be implemented in FICON VSANs. You can still perform fabric binding configuration in a non-FICON VSAN—these configurations will only come into effect after FICON is enabled.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol in FICON networks to ensure that the list of authorized switches is identical in all switches in the fabric.

This section contains the following topics:

- [Port Security versus Fabric Binding, page 21-33](#)
- [Enforcing Fabric Binding, page 21-34](#)
- [Enabling Fabric Binding, page 21-34](#)
- [Configuring a List of sWWNs, page 21-35](#)
- [Activating Fabric Binding, page 21-35](#)
- [Saving Fabric Binding Configurations, page 21-36](#)
- [Clearing the Fabric Binding Statistics, page 21-37](#)
- [Deleting the Fabric Binding Database, page 21-37](#)
- [Verifying Fabric Binding Configurations, page 21-37](#)

Port Security versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other (see [Table 21-2](#)).

Table 21-2 Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Configured using a set of sWWN and a persistent Domain ID.	Configured using pWWNs/nWWNs or fWWNs/switch WWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.
Only the configured sWWN stored in the fabric binding database will be authorized to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port(s). The switchport, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By Binding these two devices, you lock these two ports into a group (list).
Activation is required on a per VSAN basis.	Activation is required on a per VSAN basis.
User defines specific switches which are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	User specifies the specific physical port(s) to which another device can connect
Does not learn logging in switches.	Learns about switches/devices if in learning mode.

Port-level Checking for xE ports

- Switch login—uses both Port Binding as well as the Fabric Binding feature for a given VSAN.
- Binding checks are done on the port VSAN:
 - E-port security binding check done on port VSAN.
 - TE-port security bindings check done in each vsan allowed.

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Enforcing Fabric Binding

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up (**no shutdown** command). However enforcement of fabric binding at the time of activation happens only if the VSAN is a FICON VSAN.

The fabric binding feature requires all sWWNs connecting to a switch and their persistent domain IDs to be part of the fabric-binding active database.

To configure fabric binding in each switch in the fabric, follow these steps.

-
- | | |
|---------------|---|
| Step 1 | Enable the fabric configuration feature (see the |
| Step 2 | Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric (see the “Configuring a List of sWWNs” section on page 21-35). |
| Step 3 | Activate the fabric binding database (see the “Activating Fabric Binding” section on page 21-35). |
| Step 4 | Save the fabric binding configuration (see the “Saving Fabric Binding Configurations” section on page 21-36). |
| Step 5 | Verify the fabric binding configuration (see the “Verifying Fabric Binding Configurations” section on page 21-37). |
-

Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participate in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# fabric-binding enable	Enables fabric binding on that switch.
	switch(config)# no fabric-binding enable	Disables (default) fabric binding on that switch.

View the status of the fabric binding feature of an fabric binding-enabled switch by issuing the **show fabric-binding status** command.

```
switch# show fabric-binding status
VSAN 1 :Activated database
VSAN 4 :No Active database
```

Configuring a List of sWWNs

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If a sWWN attempts to join the fabric, and that sWWN is not in the list or the sWWN is using a domain ID that differs from the one specified in the allowed list, the ISL between the switch and the fabric is automatically isolated in that VSAN and the switch is denied entry into the fabric.

The persistent domain ID must be specified along with the sWWN. Domain ID authorization is required in FICON VSANs where the domains are statically configured and the end devices reject a domain ID change in all switches in the fabric.

To configure a list of sWWNs and domain IDs, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fabric-binding database vsan 5 switch(config-fabric-binding)#	Enters the fabric binding submode for the specified VSAN.
	switch(config)# no fabric-binding database vsan 10	Deletes the fabric binding database for the specified VSAN.
Step 3	switch(config-fabric-binding)# swwn 21:00:05:30:23:11:11 domain 102	Adds the sWWN and domain ID of a switch to the configured database list.
Step 4	switch(config-fabric-binding)# swwn 21:00:05:30:23:1a:11:03 domain 101	Adds the sWWN and domain ID of another switch to the configured database list.
Step 5	switch(config-fabric-binding)# no swwn 21:00:15:30:23:1a:11:03 domain 101	Deletes the sWWN and domain ID of a switch from the configured database list.
Step 6	switch(config-fabric-binding)# exit switch(config)#	Exits the fabric binding submode.

Activating Fabric Binding

The fabric binding maintains a configuration database (config-database) and an active database. The config-database is a read-write database which collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config-database. The active database is read-only and is the database that checks each switch that attempts to login.

By default, the fabric binding feature is not activated. You cannot activate the switch if entries existing in the config database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config database.



Note

You can choose the **force** option to override these situations.

To activate the fabric binding feature, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fabric-binding activate vsan 1	Activates the fabric binding database for the specified VSAN.
	switch(config)# no fabric-binding activate vsan 10	Deactivates the fabric binding database for the specified VSAN.



Note

After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login due to fabric binding restrictions will be reinitialized.

Forcing Fabric Binding Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the **force** option.

To forcefully activate the fabric binding database, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fabric-binding activate vsan 3 force	Activates the fabric binding database for the specified VSAN forcefully—even if the configuration is not acceptable.
	switch(config)# no fabric-binding activate vsan 1 force	Reverts to the previously-configured state or to the factory default (if no state is configured).

Saving Fabric Binding Configurations

When you save the fabric binding configuration, the config-database and the active database are both saved to the startup configuration and will be available after a reboot.

- Use the **fabric-binding database copy vsan** command to copy from the active database to the configuration database. If the configured database is empty, this command is not accepted.
switch# **fabric-binding database copy vsan 1**
- Use the **fabric-binding database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.
switch# **fabric-binding database diff active vsan 1**
- Use the **fabric-binding database diff config vsan** command to obtain information on the differences between the configuration database and the active database.
switch# **fabric-binding database diff config vsan 1**

**Caution**

You cannot deactivate or disable fabric binding in a FICON-enabled VSAN.

Clearing the Fabric Binding Statistics

Use the **clear fabric-binding statistics** command to clear all existing statistics from the fabric binding database for a specified VSAN.

```
switch# clear fabric-binding statistics vsan 1
```

Deleting the Fabric Binding Database

Use the **no fabric-binding** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no fabric-binding database vsan 1
```

Verifying Fabric Binding Configurations

Use the **show** commands to display all fabric binding information configured on this switch (see Examples 21-14 to 21-22).

Example 21-14 Displays Configured Fabric Binding Database Information

```
switch# show fabric-binding database
```

```
-----
Vsan   Logging-in Switch WWN      Domain-id
-----
1       21:00:05:30:23:11:11:11    0x66 (102)
1       21:00:05:30:23:1a:11:03    0x19 (25)
1       20:00:00:05:30:00:2a:1e    0xea (234)
4       21:00:05:30:23:11:11:11    0x66 (102)
4       21:00:05:30:23:1a:11:03    0x19 (25)
61      21:00:05:30:23:1a:11:03    0x19 (25)
61      21:00:05:30:23:11:11:11    0x66 (102)
[Total 7 entries]
```

Example 21-15 Displays Active Fabric Binding Information

```
switch# show fabric-binding database active
```

```
-----
Vsan   Logging-in Switch WWN      Domain-id
-----
1       21:00:05:30:23:11:11:11    0x66 (102)
1       21:00:05:30:23:1a:11:03    0x19 (25)
1       20:00:00:05:30:00:2a:1e    0xea (234)
61      21:00:05:30:23:1a:11:03    0x19 (25)
61      21:00:05:30:23:11:11:11    0x66 (102)
61      20:00:00:05:30:00:2a:1e    0xef (239)
```

Example 21-16 Displays Active VSAN-Specific Fabric Binding Information

```
switch# show fabric-binding database active vsan 61
-----
Vsan   Logging-in Switch WWN      Domain-id
-----
61      21:00:05:30:23:1a:11:03      0x19(25)
61      21:00:05:30:23:11:11:11      0x66(102)
61      20:00:00:05:30:00:2a:1e      0xef(239)
[Total 3 entries]
```

Example 21-17 Displays Configured VSAN-Specific Fabric Binding Information

```
switch# show fabric-binding database vsan 4
-----
Vsan   Logging-in Switch WWN      Domain-id
-----
4       21:00:05:30:23:11:11:11      0x66(102)
4       21:00:05:30:23:1a:11:03      0x19(25)
[Total 2 entries]
```

Example 21-18 Displays Fabric Binding Statistics

```
switch# show fabric-binding statistics
Statistics For VSAN: 1
-----
Number of sWWN permit: 0
Number of sWWN deny   : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 4
-----
Number of sWWN permit: 0
Number of sWWN deny   : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 61
-----
Number of sWWN permit: 0
Number of sWWN deny   : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 345
-----
Number of sWWN permit: 0
Number of sWWN deny   : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 346
-----
Number of sWWN permit: 0
Number of sWWN deny   : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 347
-----
Number of sWWN permit: 0
```

```

Number of sWWN deny : 0

Total Logins permitted : 0
Total Logins denied : 0
Statistics For VSAN: 348
-----
Number of sWWN permit: 0
Number of sWWN deny : 0

Total Logins permitted : 0
Total Logins denied : 0
Statistics For VSAN: 789
-----
Number of sWWN permit: 0
Number of sWWN deny : 0

Total Logins permitted : 0
Total Logins denied : 0
Statistics For VSAN: 790
-----
Number of sWWN permit: 0
Number of sWWN deny : 0

Total Logins permitted : 0
Total Logins denied : 0

```

Example 21-19 Displays Fabric Binding Status for Each VSAN

```

switch# show fabric-binding status
VSAN 1 :Activated database
VSAN 4 :No Active database
VSAN 61 :Activated database
VSAN 345 :No Active database
VSAN 346 :No Active database
VSAN 347 :No Active database
VSAN 348 :No Active database
VSAN 789 :No Active database
VSAN 790 :No Active database

```

Example 21-20 Displays Fabric Binding Violations

```

switch# show fabric-binding violations
-----
VSAN Switch WWN [domain] Last-Time [Repeat count] Reason
-----
3 20:00:00:05:30:00:4a:1e [*] Nov 25 05:44:58 2003 [2] sWWN not found
3 20:00:00:05:30:00:4a:1e [0xeb] Nov 25 05:46:14 2003 [2] Domain mismatch
4 20:00:00:05:30:00:4a:1e [*] Nov 25 05:46:25 2003 [1] Database mismatch

```



Note

Note: In VSAN 100, the * indicates that the sWWN itself was not found in the list. In VSAN 2, the sWWN was found in the list, but has a domain ID mismatch.

Example 21-21 Displays EFMD Statistics

```

switch# show fabric-binding efmd statistics

EFMD Protocol Statistics for VSAN 1
-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts -> Transmitted : 0 , Received : 0

```

```

Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

```

EFMD Protocol Statistics for VSAN 4

```

-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

```

EFMD Protocol Statistics for VSAN 61

```

-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

```

Example 21-22 Displays EFMD Statistics for a Specified VSAN

```
switch# show fabric-binding efmd statistics vsan 4
```

EFMD Protocol Statistics for VSAN 4

```

-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

```


Displaying RLIR Information

The Registered Link Incident Report (RLIR) application provides a method for a switchport to send a LIR to a registered Nx-port.

When a Link Incident Record (LIR) is detected in FICON-enabled switches in the Cisco MDS 9000 Family form a RLIR Extended Link Service (ELS) and sends it to the members in it's Established Registration List (ERL).

In case of multi-switch topology, a Distribute Registered Link Incident Record (DRLIR) Inter Link Service (ILS) are sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends to the members of the ERL.

The Nx-ports interested in receiving the RLIR ELS send Link Incident Record Registration (LIRR) ELS request to the management server on the Switch. The RLIRs are processed on a per-VSAN basis.

The RLIR application is highly available and the data is stored to persistent storage when the **copy running-config startup-config** command is issued.

The **show rlir statistics** command displays the complete statistics of LIRR, RLIR, and DRLIR frames. It lists the number of frames received, sent, and rejected. Specify the VSAN ID for per VSAN statistics, If you do not specify the VSAN ID, then the statistics is shown for all active VSANs (see Examples 21-23 and 21-24).

Example 21-23 Displays RLIR Statistics for All VSANs

```
switch# show rlir statistics

Statistics for VSAN: 1
-----

Number of LIRR received      = 0
Number of LIRR ACC sent     = 0
Number of LIRR RJT sent     = 0
Number of RLIR sent         = 0
Number of RLIR ACC received = 0
Number of RLIR RJT received = 0
Number of DRLIR received    = 0
Number of DRLIR ACC sent    = 0
Number of DRLIR RJT sent    = 0
Number of DRLIR sent        = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0

Statistics for VSAN: 100
-----

Number of LIRR received      = 26
Number of LIRR ACC sent     = 26
Number of LIRR RJT sent     = 0
Number of RLIR sent         = 815
Number of RLIR ACC received = 815
Number of RLIR RJT received = 0
Number of DRLIR received    = 417
Number of DRLIR ACC sent    = 417
Number of DRLIR RJT sent    = 0
Number of DRLIR sent        = 914
Number of DRLIR ACC received = 828
Number of DRLIR RJT received = 0
```

Example 21-24 Displays RLIR Statistics for a Specified VSAN

```
switch# show rlir statistics vsan 4
```

```
Statistics for VSAN: 4
-----

Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
Number of DRLIR received    = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0
```

The **show rlir erl** command shows the list of Nx-ports that are registered to receive the RLIRs with the switch. If the VSAN ID is not specified, the details are shown for all active VSANs (see Examples 21-25 and 21-26).

Example 21-25 Displays All ERLs

```
switch# show rlir erl
```

```
Established Registration List for VSAN: 2
-----
FC-ID          LIRR FORMAT    REGISTERED FOR
-----
0x0b0200       0x18           always receive
Total number of entries = 1

Established Registration List for VSAN: 100
-----
FC-ID          LIRR FORMAT    REGISTERED FOR
-----
0x0b0500       0x18           conditional receive
0x0b0600       0x18           conditional receive
Total number of entries = 2
```

In [Example 21-25](#), if the Registered For column states that an FC ID is conditional receive, the source port is registered as a valid recipient of subsequent RLIRs. This source port is selected as an RLIR recipient only if no other ERL recipient is selected.

In [Example 21-25](#), if the Registered For column states that an FC ID is always receive, the source port is registered as valid recipient of subsequent RLIRs. This source port is always selected as a LIR recipient.



Note

If an *always receive* RLIR is not registered for any N-port or if the delivery of an RLIR fails for one of those ports, then the RLIR is sent to a port registered to *conditional receive* RLIRs.

Example 21-26 Displays ERLs for the Specified VSAN

```
switch# show rlir erl vsan 100
Established Registration List for VSAN: 100
-----
FC-ID          LIRR FORMAT    REGISTERED FOR
-----
0x0b0500       0x18           conditional receive
0x0b0600       0x18           conditional receive

Total number of entries = 2
```

**Note**

In Examples 21-27, 21-28, and 21-29, if the host timestamp (marked by the *) is available, it is printed along with the switch timestamp. If the host timestamp is not available, only the switch timestamp is printed.

Example 21-27 Displays the LIR History

```
switch# show rlir history

Link incident history
-----
*Host Time Stamp
Switch Time Stamp          Port    Interface    Link Incident
-----
*Sun Nov 30 21:47:28 2003
Sun Nov 30 13:47:55 2003      2      fc1/2      Implicit Incident
*Sun Nov 30 22:00:47 2003
Sun Nov 30 14:01:14 2003      2      fc1/2      NOS Received
*Sun Nov 30 22:00:55 2003
Sun Nov 30 14:01:22 2003      2      fc1/2      Implicit Incident
*Mon Dec 1 20:14:26 2003
Mon Dec 1 12:14:53 2003      4      fc1/4      Implicit Incident
*Mon Dec 1 20:14:26 2003
Mon Dec 1 12:14:53 2003      4      fc1/4      Implicit Incident
*Thu Dec 4 04:43:32 2003
Wed Dec 3 20:43:59 2003      2      fc1/2      NOS Received
*Thu Dec 4 04:43:41 2003
Wed Dec 3 20:44:08 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 04:46:53 2003
Wed Dec 3 20:47:20 2003      2      fc1/2      NOS Received
*Thu Dec 4 04:47:05 2003
Wed Dec 3 20:47:32 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 04:48:07 2003
Wed Dec 3 20:48:34 2003      2      fc1/2      NOS Received
*Thu Dec 4 04:48:39 2003
Wed Dec 3 20:49:06 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 05:02:20 2003
Wed Dec 3 21:02:47 2003      2      fc1/2      NOS Received
*Thu Dec 4 05:02:29 2003
Wed Dec 3 21:02:56 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 05:02:47 2003
Wed Dec 3 21:03:14 2003      4      fc1/4      NOS Received
*Thu Dec 4 05:02:54 2003
Wed Dec 3 21:03:21 2003      4      fc1/4      Implicit Incident
*Thu Dec 4 05:02:54 2003
Wed Dec 3 21:03:21 2003      4      fc1/4      Implicit Incident
...
```

Example 21-28 Displays Recent LRIs for a Specified Interface

```
switch# show rlir recent interface fc1/1-16
Recent link incident records
-----
*Host Time Stamp
Switch Time Stamp      Port    Interface  Link Incident
-----
*Thu Dec 4 05:02:29 2003
Wed Dec 3 21:02:56 2003    2        fc1/2      Implicit Incident
*Thu Dec 4 05:02:54 2003
Wed Dec 3 21:03:21 2003    4        fc1/4      Implicit Incident
```

Example 21-29 Displays Recent LRIs for a Specified Port Number

```
switch# show rlir recent portnumber 1-16
Recent link incident records
-----
*Host Time Stamp
Switch Time Stamp      Port    Interface  Link Incident
-----
*Thu Dec 4 05:02:29 2003
Wed Dec 3 21:02:56 2003    2        fc1/2      Implicit Incident
*Thu Dec 4 05:02:54 2003
Wed Dec 3 21:03:21 2003    4        fc1/4      Implicit Incident
```

Clearing RLIR Information

Use the **clear rlir statistics** command to clear all existing statistics for a specified VSAN.

```
switch# clear rlir statistics vsan 1
```

Use the **clear rlir history** command to clear the RLIR history where all link incident records are logged for all interfaces.

```
switch# clear rlir history
```

Use the **clear rlir recent interface** command to clear the most recent RLIR information for a specified interface.

```
switch# clear rlir recent interface fc 1/2
```

Use the **clear rlir recent portnumber** command to clear the most recent RLIR information for a specified port number.

```
switch# clear rlir recent portnumber 16
```



Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services modules extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and allows IP hosts to access Fibre Channel storage using iSCSI protocol.

This chapter includes the following sections:

- [IP Storage Services Module, page 22-2](#)
- [Configuring Gigabit Ethernet Interfaces, page 22-4](#)
- [Configuring FCIP, page 22-17](#)
- [Configuring iSCSI, page 22-43](#)
- [Configuring Storage Name Services, page 22-91](#)
- [Default IP Storage Settings, page 22-95](#)



Note

FCIP and iSCSI features are specific to the IPS module and can be implemented in Cisco MDS 9216 switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 1.1(x) or above.

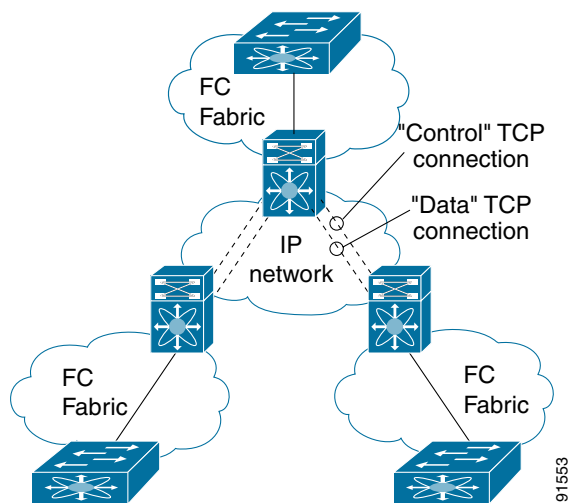
IP Storage Services Module

The IPS services module (IPS module) allows you to use FCIP and iSCSI features. It integrates seamlessly into the Cisco MDS 9000 Family, and supports the full range of features available on other switching modules, including VSANs, security, and traffic management.

The IPS module can be used in any Cisco MDS 9000 Family switch and has eight Gigabit Ethernet ports. Each port can run FCIP and iSCSI protocols simultaneously.

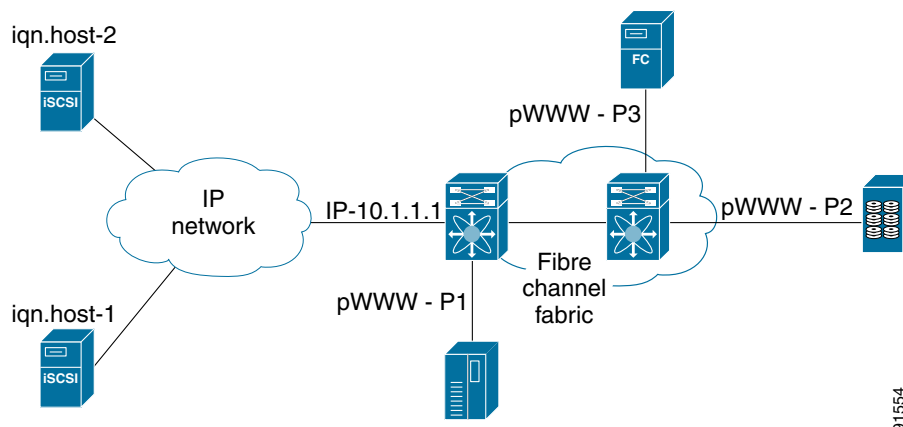
- **FCIP**—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. [Figure 22-1](#) depicts the FCIP scenarios in which the IPS module is used.

Figure 22-1 FCIP Scenarios



- **iSCSI**—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a MDS 9000 IPS port over a Transmission Control Protocol (TCP)/Internet Protocol (IP) connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. [Figure 22-2](#) depicts the iSCSI scenarios in which the IPS module is used.

Figure 22-2 iSCSI Scenarios



Verifying the Module Status

After inserting the module, verify the status of the module using the **show module** command:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  -
2    16     1/2 Gbps FC Module        DS-X9016            ok
4    8      IP Storage Module        DS-X9308-SMIP      ok <-----IPS module
5    0      Supervisor/Fabric-1       DS-X9530-SF1-K9     active *
6    0      Supervisor/Fabric-1       DS-X9530-SF1-K9     ha-standby

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  -
2    1.1(1)      0.3         20:41:00:05:30:00:86:5e to 20:50:00:05:30:00:86:5e
4    1.1(1)      0.2         20:c1:00:05:30:00:86:5e to 20:c8:00:05:30:00:86:5e
5    1.1(1)      0.602      --
6    1.1(1)      0.602      --

Mod  MAC-Address(es)                Serial-Num
---  -
2    00-05-30-00-9f-62 to 00-05-30-00-9f-66  JAB064505YV
4    00-05-30-00-a1-ae to 00-05-30-00-a1-ba  JAB0649059h
5    00-05-30-00-9f-f6 to 00-05-30-00-9f-fa  JAB06350B1M
6    00-05-30-00-9f-f2 to 00-05-30-00-9f-f6  JAB06350B1F

* this terminal session
```

Configuring Gigabit Ethernet Interfaces

This section includes the following topics:

- [About Gigabit Ethernet Interfaces, page 22-4](#)
- [Basic Gigabit Ethernet Configuration, page 22-4](#)
- [About VLANs for Gigabit Ethernet, page 22-5](#)
- [VLAN Configuration, page 22-6](#)
- [Interface Subnet Requirements, page 22-6](#)
- [Managing IP Routing, page 22-7](#)
- [Verifying Gigabit Ethernet Connectivity, page 22-7](#)
- [Managing ARP Caches, page 22-8](#)
- [Displaying Statistics, page 22-8](#)
- [Gigabit Ethernet High Availability, page 22-12](#)
- [Configuring CDP, page 22-16](#)
- [IPS Core Dumps, page 22-16](#)

About Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On the IPS module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

A new port mode, called **IPS**, is defined for Gigabit Ethernet ports on the IPS module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.



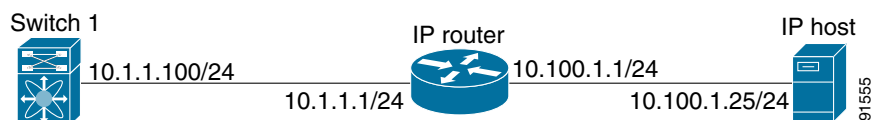
Tip

Gigabit Ethernet ports on the IPS module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Basic Gigabit Ethernet Configuration

Figure 22-3 depicts a basic Gigabit Ethernet configuration.

Figure 22-3 Gigabit Ethernet Configuration



To configure the Gigabit Ethernet interface for the scenario in [Figure 22-3](#), follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.100 255.255.255.0	Enters the IP address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

You can configure the switch to receive and transfer large (or jumbo) frames on a port. The default IP MTU frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased to 9000 bytes. The following example sets the size to 3000 bytes. Independent of the MTU size, the IPS module does not pack multiple IP frames (converted to FCIP or to iSCSI).



Note

The minimum MTU size for a port running iSCSI is 512 bytes.

To configure MTU frame size, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot2, port 2).
Step 3	switch(config-if)# switchport mtu 3000	Changes the IP maximum transmission unit (MTU) to 3000. The default is 1500.

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

IPS gigabit ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one IPS port, configure subinterfaces—one for each VLAN. Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name (the slot-number/><port-number>.<VLAN-ID>).



Note

If the IPS module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- the Ethernet switch port connected to the IPS module is configured as a trunking port, and
- the encapsulation is set to 802.1Q and not ISL, which is the default.

VLAN Configuration

To configure a VLAN subinterface (the VLAN ID), follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2.100 switch(config-if)#	Specifies the subinterface on which 802.1Q is used (slot2, port 2, VLAN ID 100). Note The subinterface number, 100 in this example, is the VLAN ID. The VLAN ID ranges from 1 to 4093.
Step 3	switch(config-if)# ip address 10.1.1.100 255.255.255.0	Enters the IP address (10.1.1.100) and IP mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the interface.

Interface Subnet Requirements

Gigabit Ethernet interface (major), subinterfaces (VLAN ID) and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 22-1](#)).

Table 22-1 Subnet Requirements for Interfaces

Interface 1	Interface 2	Same Subnet Allowed	Notes
Gigabit Ethernet 1/1	Gigabit Ethernet 1/2	Yes	Two major interfaces can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.100	Yes	Two subinterfaces with the same VLAN ID can be configured in the same or different subnets.
Gigabit Ethernet 1/1.100	Gigabit Ethernet 1/2.200	No	Two subinterfaces with different VLAN ID cannot be configured in the same subnet.
Gigabit Ethernet 1/1	Gigabit Ethernet 1/1.100	No	A VLAN ID cannot be configured on the same subnet as the major interface.
mgmt0	Gigabit Ethernet 1/1.100	No	The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces.
mgmt0	Gigabit Ethernet 1/1	No	



Note

The configuration requirements in [Table 22-1](#) also applies to Ethernet PortChannels.

Managing IP Routing

To configure static IP routing (see [Figure 22-3](#)) through the Gigabit Ethernet interface, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# ip route 10.100.1.0 255.255.255.0 10.1.1.1 switch(config-if)#	Enters the IP subnet (10.100.1.0 255.255.255.0) of the IP host and configures the next hop 10.1.1.1, which is the IP address of the router connected to the Gigabit Ethernet interface.

Displaying the IP Route Table

The **show ips ip route interface ethernet** command takes the ethernet interface as a parameter and returns the route table for the interface. See [Example 22-1](#).

Example 22-1 Displays the Route Table

```
switch# show ips ip route interface gig 8/1
Codes: C - connected, S - static
No default gateway
C 10.1.3.0/24 is directly connected, GigabitEthernet8/1
```

Connected (C) identifies the subnet in which the interface is configured (directly connected to the interface). Static (S) identifies the static routes that go through the router.

Verifying Gigabit Ethernet Connectivity

The **ping** command sends echo request packets out to a remote device at an IP address that you specify (see the “[Using the ping Command](#)” section on page 2-14).

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch using the **ping** command. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly. See [Example 22-2](#).

Example 22-2 Verifying Gigabit Ethernet Connectivity

```
switch# ping 10.100.1.25
PING 10.100.1.25 (10.100.1.25): 56 data bytes
64 bytes from 10.100.1.25: icmp_seq=0 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=1 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=2 ttl=255 time=0.1 ms
--- 10.100.1.25 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.1 ms
```



Note

If the connection fails, verify the following, and repeat the **ping** command:

- the IP address for the destination (IP host) is correctly configured,
- the host is active (powered on),

- the IP route is configured correctly,
- the IP host has a route to get to the Gigabit Ethernet interface subnet, and
- the Gigabit Ethernet interface is in the up state (use the **show interface gigabitethernet** command).

Managing ARP Caches

Use the **show ips arp interface gigabitethernet** command to display the ARP cache on the Gigabit Ethernet interfaces. This command takes the Ethernet interface as a parameter and returns the ARP cache for that interface. See [Example 22-3](#).

Example 22-3 Displays ARP Caches

```
switch# show ips arp interface gigabitethernet 7/1
Protocol      Address      Age (min)    Hardware Addr  Type   Interface
Internet      20.1.1.5     3            0005.3000.9db6 ARPA    GigabitEthernet7/1
Internet      20.1.1.10    7            0004.76eb.2ff5 ARPA    GigabitEthernet7/1
Internet      20.1.1.11    16           0003.47ad.21c4 ARPA    GigabitEthernet7/1
Internet      20.1.1.12    6            0003.4723.c4a6 ARPA    GigabitEthernet7/1
Internet      20.1.1.13    13           0004.76f0.ef81 ARPA    GigabitEthernet7/1
Internet      20.1.1.14    0            0004.76e0.2f68 ARPA    GigabitEthernet7/1
Internet      20.1.1.15    6            0003.47b2.494b ARPA    GigabitEthernet7/1
Internet      20.1.1.17    2            0003.479a.b7a3 ARPA    GigabitEthernet7/1
...
```

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache. See Examples [22-4](#) and [22-5](#).

Example 22-4 Clearing One ARP Cache Entry

```
switch# clear ips arp address 10.2.2.2 interface gigabitethernet 8/7
arp clear successful
```

Example 22-5 Clearing All ARP Cache Entries

```
switch# clear ips arp interface gigabitethernet 8/7
arp clear successful
```



Note

Use the physical interface, not the subinterface, for all ARP cache commands.

Displaying Statistics

This section provides examples to verify Gigabit Ethernet and TCP/IP statistics on the IP storage ports.

Displaying Gigabit Ethernet Interface Statistics

Use the **show interface Gigabit Ethernet** command on each switch to verify that the interfaces are up and functioning as desired. See [Example 22-6](#).

Example 22-6 Displays the Gigabit Ethernet Interface

```
switch# show interface gigabitethernet 8/1
```

```
GigabitEthernet8/1 is up          <-----The interface is in the up state.
Hardware is GigabitEthernet, address is 0005.3000.a98e
Internet address is 10.1.3.1/24
MTU 1500 bytes, BW 1000000 Kbit
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
5 minutes input rate 744 bits/sec, 93 bytes/sec, 1 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
3343 packets input, 406582 bytes
  0 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
8 packets output, 336 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors
```

Example 22-7 Displays the Gigabit Ethernet's Subinterface

```
switch# show interface gigabitethernet 4/2.100
GigabitEthernet4/2.100 is up
Hardware is GigabitEthernet, address is 0005.3000.abcb
Internet address is 10.1.2.100/24
MTU 1500 bytes
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
0 packets input, 0 bytes
  0 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
1 packets output, 46 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors
```

Displaying Ethernet MAC Statistics

The **show ips stats mac interface gigabitethernet** command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 22-8](#).

Example 22-8 Displays Ethernet MAC Statistics

```
switch# show ips stats mac interface gigabitethernet 8/1
Ethernet MAC statistics for port GigabitEthernet8/1
Hardware Transmit Counters
  237 frame 43564 bytes
  0 collisions, 0 late collisions, 0 excess collisions
  0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
Hardware Receive Counters
  427916 bytes, 3464 frames, 0 multicasts, 3275 broadcasts
  0 bad, 0 runt, 0 CRC error, 0 length error
  0 code error, 0 align error, 0 oversize error
Software Counters
  3429 received frames, 237 transmit frames
  0 frames soft queued, 0 current queue, 0 max queue
  0 dropped, 0 low memory
```



Note

Use the physical interface, not the subinterface, to display Ethernet MAC statistics.

Displaying DMA-Bridge Statistics

You can display direct memory access (DMA) device statistics using the **show ips stats dma-bridge interface gigabitethernet** command. This command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 22-9](#).

Example 22-9 Displays DMA-Bridge Statistics

```
switch# show ips stats dma-bridge interface gigabitethernet 7/1
Dma-bridge ASIC Statistics for port GigabitEthernet7/1
  Hardware Egress Counters
    231117 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
  Hardware Ingress Counters
    218255 Good, 0 protocol error, 0 header checksum error
    0 FC CRC error, 0 iSCSI CRC error, 0 parity error
  Software Egress Counters
    231117 good frames, 0 bad header cksum, 0 bad FIFO SOP
    0 parity error, 0 FC CRC error, 0 timestamp expired error
    0 unregistered port index, 0 unknown internal type
    0 RDL ok, 0 RDL drop (too big), 0 RDL ttl_1
    3656368645 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
    Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
  Software Ingress Counters
    218255 Good frames, 0 header cksum error, 0 FC CRC error
    0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
    0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
    0 out of memory drop, 0 queue full drop
    0 RDL ok, 0 RDL drop (too big)
    Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
```

This output shows all Fibre Channel frames that ingress or egress from the Gigabit Ethernet port.



Note

Use the physical interface, not the subinterface, to display DMA-bridge statistics.

Displaying TCP/IP Statistics



Note

Use the physical interface, not the subinterface, to display TCP/IP statistics.

Use the **show ips stats ip interface gigabitethernet** to display and verify IP statistics. This command takes the main Gigabit Ethernet interface as a parameter and returns IP statistics for that interface. See [Example 22-10](#).

Example 22-10 Displays IP Statistics

```
switch# show ips stats ip interface gigabitethernet 4/1
Internet Protocol Statistics for port GigabitEthernet4/1
  168 total received, 168 good, 0 error
  0 reassembly required, 0 reassembled ok, 0 dropped after timeout
  371 packets sent, 0 outgoing dropped, 0 dropped no route
  0 fragments created, 0 cannot fragment
```

Use the **show ips stats tcp interface gigabitethernet** to display and verify TCP statistics. This command takes the main ethernet interface as a parameter, and shows TCP stats along with the connection list and TCP state. The **detail** option shows all information maintained by the interface. See Examples 22-11 and 22-12.

Example 22-11 Displays TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1
TCP Statistics for port GigabitEthernet4/1
Connection Stats
  0 active openings, 3 accepts
  0 failed attempts, 12 reset received, 3 established
Segment stats
  163 received, 355 sent, 0 retransmitted
  0 bad segments received, 0 reset sent

TCP Active Connections
  Local Address      Remote Address      State      Send-Q    Recv-Q
  0.0.0.0:3260       0.0.0.0:0           LISTEN     0          0
```

Example 22-12 Displays Detailed TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1 detail
TCP Statistics for port GigabitEthernet4/1
TCP send stats
  355 segments, 37760 bytes
  222 data, 130 ack only packets
  3 control (SYN/FIN/RST), 0 probes, 0 window updates
  0 segments retransmitted, 0 bytes
  0 retransmitted while on ethernet send queue, 0 packets split
  0 delayed acks sent
TCP receive stats
  163 segments, 114 data packets in sequence, 6512 bytes in sequence
  0 predicted ack, 10 predicted data
  0 bad checksum, 0 multi/broadcast, 0 bad offset
  0 no memory drops, 0 short segments
  0 duplicate bytes, 0 duplicate packets
  0 partial duplicate bytes, 0 partial duplicate packets
  0 out-of-order bytes, 1 out-of-order packets
  0 packet after window, 0 bytes after window
  0 packets after close
  121 acks, 37764 ack bytes, 0 ack toomuch, 4 duplicate acks
  0 ack packets left of snd_una, 0 non-4 byte aligned packets
  8 window updates, 0 window probe
  30 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops
TCP Connection Stats
  0 attempts, 3 accepts, 3 established
  3 closed, 2 drops, 0 conn drops
  0 drop in retransmit timeout, 1 drop in keepalive timeout
  0 drop in persist drops, 0 connections drained
TCP Miscellaneous Stats
  115 segments timed, 121 rtt updated
  0 retransmit timeout, 0 persist timeout
  12 keepalive timeout, 11 keepalive probes
TCP SACK Stats
  0 recovery episodes, 0 data packets, 0 data bytes
  0 data packets retransmitted, 0 data bytes retransmitted
  0 connections closed, 0 retransmit timeouts
TCP SYN Cache Stats
  15 entries, 3 connections completed, 0 entries timed out
  0 dropped due to overflow, 12 dropped due to RST
  0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
```

```

0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
0 hash collisions, 0 retransmitted
TCP Active Connections
  Local Address          Remote Address          State      Send-Q    Recv-Q
  0.0.0.0:3260           0.0.0.0:0              LISTEN     0         0

```

Use the **show ips stats icmp interface gigabitethernet** to display and verify IP statistics. This command takes the main ethernet interface as a parameter and returns the ICMP statistics for that interface. See [Example 22-13](#).

Example 22-13 Displays ICMP Statistics

```

switch# show ips stats icmp interface gigabitethernet 4/1
ICMP Statistics for port GigabitEthernet4/1
  5 ICMP messages received
  0 ICMP messages dropped due to errors
  ICMP input histogram
    5 echo request
  ICMP output histogram
    5 echo reply

```

Gigabit Ethernet High Availability

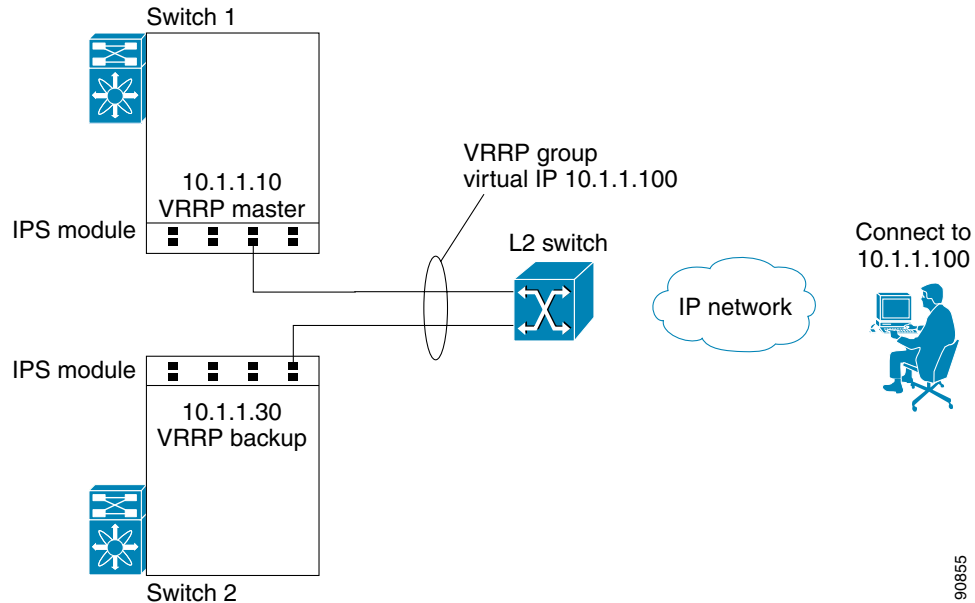
Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

Configuring VRRP

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services (see the [“Configuring VRRP” section on page 20-18](#)).

VRRP provides IP address fail over protection to an alternate Gigabit Ethernet interface so the IP address is always available (see [Figure 22-4](#)).

Figure 22-4 VRRP Scenario



In Figure 22-4, all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module
- Interfaces across IPS modules in one switch
- Interfaces across IPS modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels
- Subinterfaces

To configure VRRP for Gigabit Ethernet interfaces, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface gigabitethernet 2/2 switch(config-if)#	Enters the interface configuration mode on the Gigabit Ethernet interface (slot2, port 2).
Step 3	switch(config-if)# ip address 10.1.1.10 255.255.255.0	Enters the IP address (10.1.1.10) and IP mask (255.255.255.0) for the Gigabit Ethernet interface.
Step 4	switch(config-if)# no shutdown	Enables the selected interface.
Step 5	switch(config-if)# vrrp 100 switch(config-if-vrrp)	Creates a VR ID 100.

	Command	Purpose
Step 6	<code>switch(config-if-vrrp)# address 10.1.1.100</code>	Configures the virtual IP address (10.1.1.100) for the selected VRRP group (identified by the VR ID). Note The virtual IP address must be in the same subnet as the IP address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IP address.
Step 7	<code>switch(config-if-vrrp)# priority 10</code>	Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master.
Step 8	<code>switch(config-if-vrrp)# no shutdown</code>	Enables the VRRP protocol on the selected interface.



Note The VRRP **preempt** option is not supported on IP storage Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

Configuring Ethernet PortChannels

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

The data traffic from one TCP connection always travels on the same physical links. An Ethernet switch connecting to the MDS Gigabit Ethernet port can implement load balancing based on its IP address, its source-destination MAC address, or its IP and port. In iSCSI scenarios if the Ethernet switch is not capable of load-balancing based on the IP address or the IP port, multiple iSCSI initiators are required to take advantage of the Ethernet PortChannel feature.

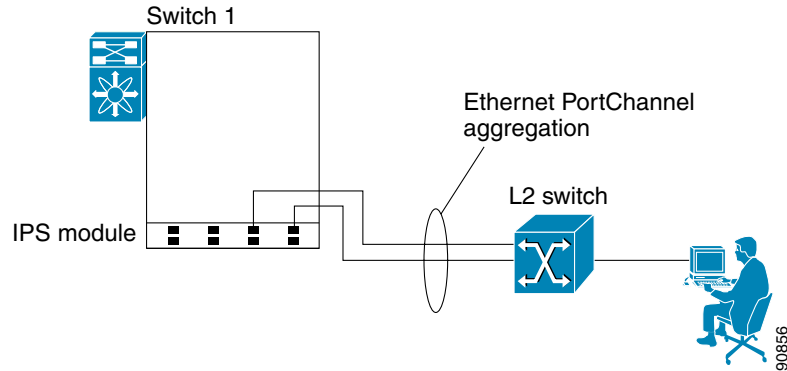


Note The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3aa protocol.

Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see [Figure 22-5](#)).



Note PortChannel members must be one of these combinations: ports 1-2, ports 3-4, ports 5-6, or ports 7-8.

Figure 22-5 Ethernet PortChannel Scenario

In [Figure 22-5](#), Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel.

**Note**

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that FCIP link.

PortChannel configuration specified in [Chapter 12, “Configuring PortChannels”](#) also apply to Ethernet PortChannel configurations.

**Note**

PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

To configure Ethernet PortChannels, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch(config)# interface port-channel 10 switch(config-if)#	Configures the specified PortChannel (10).
Step 3	switch(config-if)# ip address 10.1.1.1 255.255.255.0	Enters the IP address (10.1.1.1) and IP mask (255.255.255.0) for the PortChannel. Note A PortChannel does not have any members when first created.
Step 4	switch(config-if)# no shutdown	Enables the interface.
Step 5	switch(config)# interface gigabitethernet 9/3 switch(config-if)#	Configures the specified Gigabit Ethernet interface (slot 9, port 3).
Step 6	switch(config-if)# channel-group 10 gigabitethernet 9/3 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do “no shutdown” at both ends to bring them up switch(config-if)#	Adds Gigabit Ethernet interfaces 9/3 to channel group 10. If channel group 10 does not exist, it is created. The port is shut down.
Step 7	switch(config-if)# no shutdown	Enables the selected interface.

	Command	Purpose
Step 8	switch(config)# interface gigabitethernet 9/4 switch(config-if)#	Configures the specified Gigabit Ethernet interface (slot 9, port 4).
Step 9	switch(config-if)# channel-group 10 gigabitethernet 9/4 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up	Adds Gigabit Ethernet interfaces 9/4 to channel group 10. The port is shut down.
Step 10	switch(config-if)# no shutdown	Enables the selected interface.

**Note**

Gigabit Ethernet interfaces cannot be added to a PortChannel if one of the following cases apply:

- if the interface already has an IP address assigned, or
- if subinterfaces are configured on that interface.

Configuring CDP

The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module. See the [“Configuring CDP” section on page 4-37](#).

IPS Core Dumps

IPS core dumps are different from the system’s kernel core dumps for other modules. When the IPS module’s operating system (OS) unexpectedly resets, it is sometimes useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Core dumps take up significant space and hence the level of what gets stored can be configured using one of the two options:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files).
- Full core dumps—Each full core dump consists of 75 parts (75 files). This dump cannot be saved on the supervisor module due to its large space requirement.

To configure IPS cores on the IPS module, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# ips core dump full ips core dump full' successfully set for module 9	Configures a dump of the full core generation for the IPS module in slot 9.
	switch(config)# ips core dump partial ips core dump partial' successfully set for module 9	Configures a dump of the partial core generation for the IPS module in slot 9.

Configuring FCIP

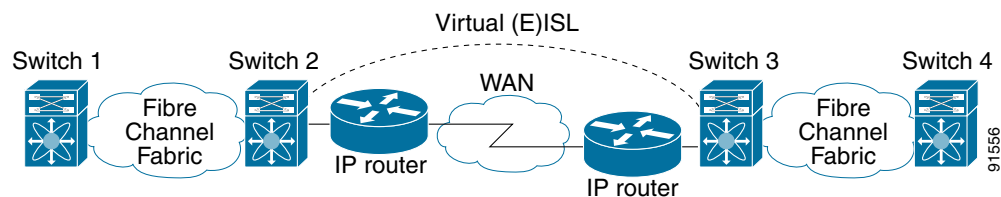
This section includes the following topics:

- [About FCIP, page 22-17](#)
- [Basic FCIP Configuration, page 22-20](#)
- [Advanced FCIP Profile Configuration, page 22-22](#)
- [Advanced FCIP Interface Configuration, page 22-28](#)
- [E Port Configurations, page 22-34](#)
- [Displaying FCIP Information, page 22-36](#)
- [FCIP High Availability, page 22-39](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 22-42](#)

About FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). See [Figure 22-6](#).

Figure 22-6 Fibre Channel SANs Connected by FCIP



FCIP uses TCP as a network layer transport.



Note

For more information about FCIP protocols, refer to the IETF standards for IP storage at <http://www.ietf.org>. Also refer to Fibre Channel standards for switch backbone connection at <http://www.t11.org> (see FC-BB-2).

To configure the IPS module for FCIP, you should have a basic understanding of the following concepts:

- [FCIP and VE Ports, page 22-18](#)
- [FCIP Link, page 22-18](#)
- [FCIP Profiles, page 22-19](#)
- [FCIP Interface, page 22-19](#)

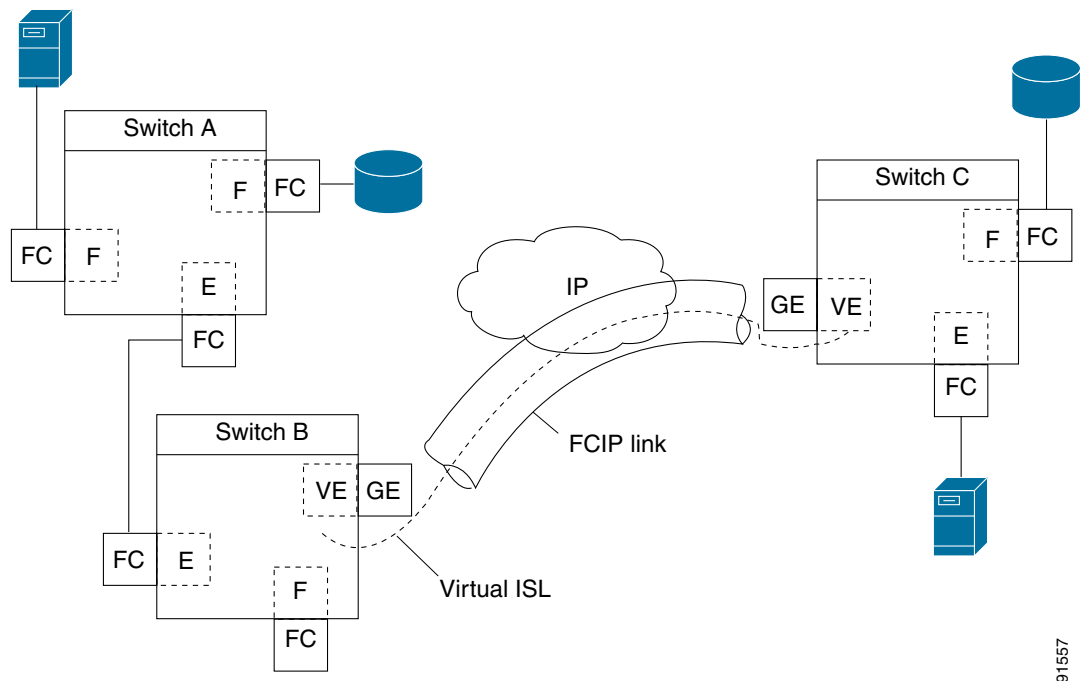
FCIP and VE Ports

Figure 22-7 describes the internal model of FCIP with respect to Fibre Channel inter switch links (ISLs) and Cisco's enhanced ISLs (EISLs). See the “E Port” section on page 10-3.

FCIP defines virtual E (VE) ports, which behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

A virtual ISL is established over a FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see Figure 22-7).

Figure 22-7 FCIP Links and Virtual ISLs



91557

FCIP Link

FCIP links consist of one or more TCP connections between two FCIP link end points. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link.

- One connection is used for data frames.
- The second connection is used only for Fibre Channel control frames, i.e. switch-to-switch protocol frames (all Class F) frames. This arrangement is used to provide low latency for all control frames.

To enable FCIP on the IPS module, a FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior is independent of the link being FCIP or pure Fibre Channel, and is based on the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port behavior is identical to E port behavior for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

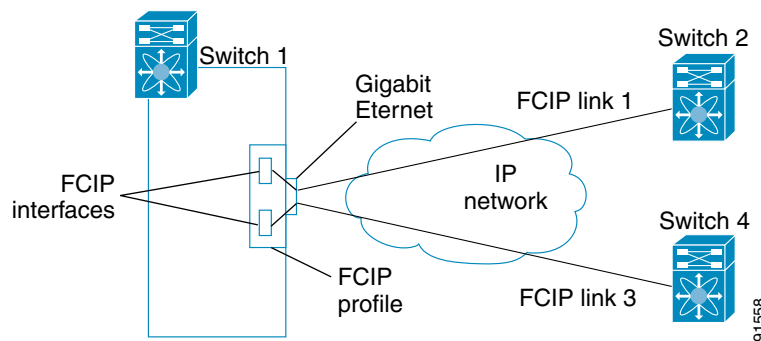
FCIP Profiles

The FCIP profile contains information about local IP address and TCP parameters. The profile defines the following information:

- the local connection points (IP address and TCP port number)
- the behavior of the underlying TCP connections for all FCIP links that use this profile

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminates (see [Figure 22-8](#)).

Figure 22-8 FCIP Profile and FCIP Links



FCIP Interface

The FCIP interface is the local end point of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.

Enabling FCIP

To begin configuring the FCIP feature, you must explicitly enable FCIP on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the FCIP feature are only available when FCIP is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To enable FCIP on any participating switch, follow these steps:

	Command	Purpose
Step 1	switch# conf t	Enters configuration mode.
Step 2	switch(config)# fcip enable	Enables FCIP on that switch.
	switch(config)# no fcip enable	Disables (default) FCIP on that switch.

Basic FCIP Configuration

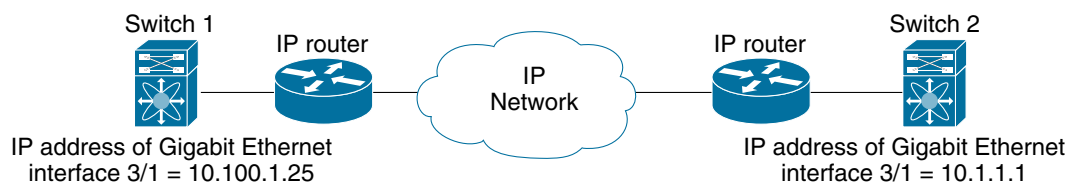
To configure a FCIP link, perform this procedure on both switches.

- Step 1** Configure the Gigabit Ethernet interface.
- Step 2** Create a FCIP profile, assign the Gigabit Ethernet interface's IP address to the profile. See the [“Creating FCIP Profiles” section on page 22-20](#).
- Step 3** Create a FCIP interface, assign the profile to the interface. See the [“Creating FCIP Links” section on page 22-21](#).
- Step 4** Configure the peer IP address for the FCIP interface. See the [“Creating FCIP Links” section on page 22-21](#).
- Step 5** Enable the interface. See the [“Creating FCIP Links” section on page 22-21](#).

Creating FCIP Profiles

To create a FCIP profile, you must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile (see [Figure 22-9](#)).

Figure 22-9 Assigning Profiles to Each Gigabit Ethernet Interface



91561

To create a FCIP profile in switch 1, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch1(config)#	Enters configuration mode.
Step 2	switch1(config)# fcip profile 10 switch1(config-profile)#	Creates a profile for the FCIP connection. The valid range is from 1 to 255.
Step 3	switch1(config-profile)# ip address 10.100.1.25	Associates the profile (10) with the local IP address of the Gigabit Ethernet interface (3/1).

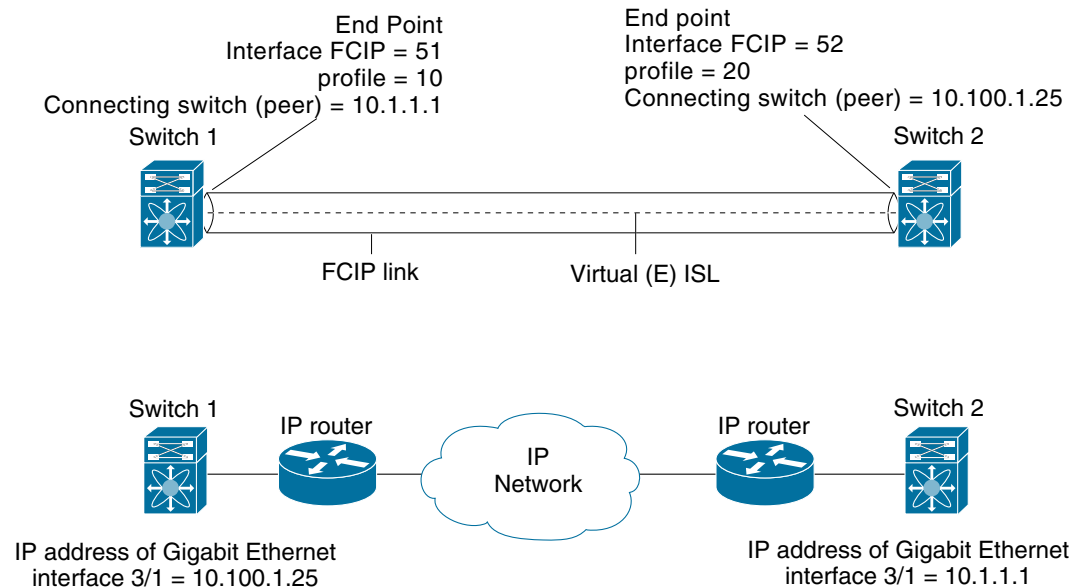
To assign FCIP profile in switch 2, follow these steps:

	Command	Purpose
Step 1	switch2# config terminal switch2(config)#	Enters configuration mode.
Step 2	switch2(config)# fcip profile 20 switch2(config-profile)#	Creates a profile for the FCIP connection.
Step 3	switch2(config-profile)# ip address 10.1.1.1	Associates the profile (20) with the local IP address of the Gigabit Ethernet interface.

Creating FCIP Links

When two FCIP link end points are created, a FCIP link is established between the two IPS modules. To create a FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) a FCIP link to that peer switch (see [Figure 22-10](#)).

Figure 22-10 Assigning Profiles to Each Gigabit Ethernet Interface



91562

To create a FCIP link end point in switch 1, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch(config)#	Enters configuration mode.
Step 2	switch1(config)# interface fcip 51 switch1(config-if)#	Creates a FCIP interface (51).
Step 3	switch1(config-if)# use-profile 10	Assigns the profile (10) to the FCIP interface.
Step 4	switch1(config-if)# peer-info ipaddr 10.1.1.1	Assigns the peer IP address information (10.1.1.1 for switch 2) to the FCIP interface
Step 5	switch1(config-if)# no shutdown	Enables the interface.

To create a FCIP link end point in switch 2, follow these steps:

	Command	Purpose
Step 1	switch2# config terminal switch2(config)#	Enters configuration mode.
Step 2	switch2(config)# interface fcip 52 switch2(config-if)#	Creates a FCIP interface (52).
Step 3	switch2(config-if)# use-profile 20	Binds the profile (20) to the FCIP interface.
Step 4	switch2(config-if)# peer-info ip address 10.100.1.25	Assigns the peer IP address information (10.100.1.25 for switch 1) to the FCIP interface
Step 5	switch2(config-if)# no shutdown	Enables the interface.

Advanced FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the local IP address and the local port, you can specify other TCP parameters as part of the FCIP profile configuration.

- [Configuring TCP Listener Ports, page 22-22](#)
- [Configuring TCP Parameters, page 22-23](#)

FCIP configuration options can be accessed from the `switch(config-profile)#` submode prompt.

To enter the `switch(config-profile)#` prompt, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcip profile 20 switch(config-profile)#	Creates the profile (if it does not already exist). The valid range is from 1 to 255.

Configuring TCP Listener Ports

The default TCP port for FCIP is 3225. You can change this port using the **port** command.

To change the default FCIP port number (3225), follow these steps:

Step 1	Command	Purpose
	<code>switch(config-profile)# port 5000</code>	Associates the profile with the local port number (5000).
	<code>switch(config-profile)# no port</code>	Reverts to the default 3225 port.

Configuring TCP Parameters

This section provides details on the TCP parameters that can be configured to control TCP behavior in a switch. The following TCP parameters can be configured.

- [Minimum Retransmit Timeout, page 22-24](#)
- [Keepalive Timeout, page 22-24](#)
- [Maximum Retransmissions, page 22-24](#)
- [Path MTU, page 22-25](#)
- [SACK, page 22-25](#)
- [Window Management, page 22-25](#)
- [Buffer Size, page 22-26](#)
- [Quality of Service, page 22-26](#)
- [Monitoring Window Congestion, page 22-27](#)

Minimum Retransmit Timeout

The **tcp minimum-retransmit-time** option controls the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds.

To configure the minimum retransmit time, follow these steps:

Step 1	Command	Purpose
	<code>switch(config-profile)# tcp min-retransmit-time 500</code>	Specifies the minimum TCP retransmit time for the TCP connection in milliseconds (500). The default is 200 milliseconds and the range is from 200 to 5000 milliseconds.
	<code>switch(config-profile)# no tcp min-retransmit-time 500</code>	Reverts the minimum TCP retransmit time to the factory default of 200 milliseconds.

Keepalive Timeout

The **tcp keepalive-timeout** option enables you to configure the interval between which the TCP connection verifies if the FCIP link is functioning. This ensures that a FCIP link failure is detected quickly even when there is no traffic.

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. This command can be used to tune the time taken to detect FCIP link failures.

The first interval during which the connection is idle is 60 seconds (default). When the connection is idle for 60 seconds, 8 keepalive probes are sent at 1-second intervals. If no response is received for these 8 probes and the connection remains idle throughout, that FCIP link is automatically closed.



Note

Only the first interval (during which the connection is idle) can be changed from the default of 60 seconds. This interval is identified using the **keepalive-timeout** option. The valid range is from 1 to 7200 seconds.

To configure the keep alive timeout, follow these steps:

Step 1	Command	Purpose
	<code>switch(config-profile)# tcp keepalive-timeout 120</code>	Specifies the keepalive timeout interval for the TCP connection in seconds (120). The default is 60 seconds. The range is from 1 to 7200 seconds.
	<code>switch(config-profile)# no tcp keepalive-timeout 120</code>	Reverts the keepalive-timeout to 60 seconds.

Maximum Retransmissions

The **tcp max-retransmissions** option specifies the maximum number of times a packet is retransmitted before TCP decides to close the connection.

To configure maximum retransmissions, follow these steps:

Step 1	Command	Purpose
	<code>switch(config-profile)# tcp max-retransmissions 6</code>	Specifies the maximum number of retransmissions (6). The default is 4 and the range is from 1 to 8 retransmissions.
	<code>switch(config-profile)# no tcp max-retransmissions 6</code>	Reverts to the default of 4 retransmissions.

Path MTU

Path MTU (PMTU) is the minimum MTU on the IP network between the two end points of the FCIP link. PMTU discovery is a mechanism by which TCP learns of the PMTU dynamically and adjusts the maximum TCP segment accordingly (RFC 1191).

By default, PMTU discovery is enabled on all switches with a default timeout of 3600 seconds. If TCP reduces the size of the max segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

To configure PMTU, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# no tcp pmtu-enable</code>	Disables PMTU discovery.
	<code>switch(config-profile)# tcp pmtu-enable</code>	Enables (default) PMTU discovery with the default value of 3600 seconds.
	<code>switch(config-profile)# tcp pmtu-enable reset-timeout 90</code>	Specifies the PMTU reset timeout to 90 seconds. The default is 3600 seconds and the range is from 60 to 3600 seconds.
	<code>switch(config-profile)# no tcp pmtu-enable reset-timeout 600</code>	Leaves the PMTU in an enabled state but changes the timeout to the default of 3600 seconds.

SACK

TCP may experience poor performance when multiple packets are lost within one window. With the limited information available from cumulative acknowledgments, a TCP sender can only learn about a single lost packet per round trip time. A selective acknowledgment (SACK) mechanism helps overcome the limitations of multiple lost packets during a TCP transmission.

The receiving TCP sends back SACK advertisements to the sender. The sender can then retransmit only the missing data segments. By default, SACK is enabled on Cisco MDS 9000 Family switches.

To configure SACK, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# no tcp sack-enable</code>	Disables SACK.
	<code>switch(config-profile)# tcp sack-enable</code>	Enables SACK (default).

Window Management

The optimal TCP window size is computed using the **max-bandwidth** option, the **min-available-bandwidth** option, and the dynamically-measured round-trip-time (RTT). The interaction and the resulting TCP behavior is outlined below:



Note

The configured **round-trip-time** option determines the window scaling factor of the TCP connection. This option is only an approximation. The measured RTT value overrides the **round-trip-time** option for window management. If the configured **round-trip-time** is too small compared to the measured RTT, then the link may not be fully utilized due to the window scaling factor being too small.

- If the average rate of the fc traffic over the preceding RTT is less than the min-available-bandwidth * RTT, every FC burst is sent immediately at the min-available-bandwidth rate, provided no TCP drops occur.

- If the average rate of the FC traffic is greater than $\text{min-available-bandwidth} * \text{RTT}$, but less than $\text{max-bandwidth} * \text{RTT}$, then if the FC traffic is transmitted in burst sizes smaller than the configured CWM value all the bursts are sent immediately by FCIP at the max-bandwidth rate.
- If the average rate of the FC traffic is larger than the $\text{min-available-bandwidth} * \text{RTT}$ and the burst size is greater than the CWM value, some traffic will not be sent immediately.

The **maximum-bandwidth** option and the measured **round-trip-time** together determine the maximum window size.

The **min-available-bandwidth** option and the measured **round-trip-time** together determine the threshold below which TCP aggressively maintains a window size sufficient to transmit at **min-available-bandwidth**. The software uses standard TCP rules to increase the window beyond the one required to maintain the **min-available-bandwidth** in order to reach the **max-bandwidth**. The defaults are **max-bandwidth** = 1G, **min-available-bandwidth** = 15 Mbps, and **round-trip-time** = 1 ms.

To configure window management, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code>	Configures the maximum available bandwidth at 900 Mbps, the minimum slow start threshold as 300 Mbps, and the round trip time as 10 milliseconds.
	<code>switch(config-profile)# no tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code>	Reverts to the factory defaults. The defaults are max-bandwidth = 1G, min-available-bandwidth = 15 Mbps and round-trip-time is 1 ms.
	<code>switch(config-profile)# tcp max-bandwidth-kbps 2000 min-available-bandwidth-kbps 2000 round-trip-time-us 200</code>	Configures the maximum available bandwidth at 2000 Kbps, the minimum slow start threshold as 2000 Kbps, and the round trip time as 200 microseconds.

Buffer Size

The **send-buffer-size** option defines the required additional buffering—beyond the normal send window size—that TCP allows before flow controlling the switch's egress path for the FCIP interface. The default buffer size is 0 KB.

To set the buffer size, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp send-buffer-size 5000</code>	Configure the advertised buffer size to 5000 KB. The valid range is from 0 to 8192 KB.
	<code>switch(config-profile)# no tcp send-buffer-size 5000</code>	Reverts the switch to its factory default (0 KB).

Quality of Service

The Quality of Service (QoS) feature specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—TOS field in the IP header).

- The control DSCP value applies to all FCIP frames in the control TCP connection.
- The data DSCP value applies to all FCIP frames in the data connection.

If the FCIP link has only one TCP connection, that data DSCP value is applied to all packets in that connection.

To set the QoS values, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# tcp qos control 3 data 5</code>	Configures the control TCP connection and data connection to mark all packets on that DSCP value.
	<code>switch(config-profile)# no tcp qos control 3 data 5</code>	Reverts the switch to its factory default (marks all packets with DSCP value 0).

Monitoring Window Congestion

The congestion window monitoring (CWM) option determines the maximum burst size allowed after an idle period.

- If the FC traffic burst is smaller than the configured CWM value, every packet is sent immediately, provided that no TCP drops were detected in the previous RTT.
- If FC traffic burst is larger than the configured CWM value, the excess packets will be sent during succeeding RTTs.

By default the **tcp cwm** option is enabled and the default burst size is 10 KB.



Tip

We recommend that this feature remain enabled to realize optimal performance. Increasing the CWM burst size can result in more packet drops in the IP network, impacting TCP performance. Only if the IP network has sufficient buffering, try increasing the CWM burst size beyond the default to achieve lower transmit latency.

To change the CWM defaults, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-profile)# no tcp cwm</code>	Disables congestion monitoring.
	<code>switch(config-profile)# tcp cwm</code>	Enables congestion monitoring and sets the default burst size at 10 KB.
	<code>switch(config-profile)# tcp cwm burstsize 30</code>	Changes the burst size to 30 KB. The valid range is from 10 to 100 KB.
	<code>switch(config-profile)# no tcp cwm burstsize 25</code>	Leaves the CWM feature in an enabled state but changes the burst size to the default of 10 KB.

Advanced FCIP Interface Configuration

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface. To do so, you must first create the interface and enter the `config-if` submode.

- [Configuring Peers, page 22-28](#)
- [Configuring Active Connection, page 22-29](#)
- [Configuring the Number of TCP Connections, page 22-30](#)
- [Enabling Time Stamps, page 22-30](#)
- [B Port Interoperability Mode, page 22-32](#)
- [Configuring FCIP Write Acceleration, page 22-35](#)

To enter the `config-if` submode, follow these steps:

	Command	Purpose
Step 1	switch# config terminal	Enters configuration mode.
Step 2	switch(config)# interface fcip 100	Creates a FCIP interface (100).

Configuring Peers

To establish a FCIP link with the peer, you can use one of two options:

- [Peer IP Address, page 22-28](#)—used to configure both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- [Special Frames, page 22-29](#)—used to configure one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the port and profile ID along with the IP address.

Peer IP Address

The basic FCIP configuration uses the peer's IP address to configure the peer information. You can also specify the peer's port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection.

To assign the peer information based on the IP address, port number, or a profile ID, follow these steps:

	Command	Purpose
Step 1	switch(config-if)# peer-info ipaddr 10.1.1.1	Assigns an IP address to configure the peer information. Since no port is specified, the default port number, 3225, is used.
	switch(config-if)# no peer-info ipaddr 10.10.1.1	Deletes the assigned peer port information.
Step 2	switch(config-if)# peer-info ipaddr 10.1.1.1 port 3000	Assigns the IP address and sets the peer TCP port to 3000. The valid port number range is from 0 to 65535.
	switch(config-if)# no peer-info ipaddr 10.1.1.1 port 2000#	Deletes the assigned peer port information.
Step 3	switch(config-if)# no shutdown	Enables the interface.

Special Frames

You can alternatively establish a FCIP link with a peer using an optional protocol called special frames. You can enable or disable the **special-frame** option. On the peer side, the **special-frame** option must be enabled in order to establish the FCIP link. When the **special-frame** option is enabled, the peer IP address (and optionally the port or the profile ID) only needs to be configured on one end of the link. Once the connection is established, a special frame is exchanged to discover and authenticate the link.

By default, the special frame feature is disabled.



Note

Refer to the Fibre Channel IP standards for further information on special frames.



Tip

Special frame negotiation provides an additional authentication security mechanism because the link validates the WWN of the peer switch.

To enable special frames, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00</code>	Enables special frames and sets the peer WWN as specified. Note The peer WWN is the WWN of the peer switch. Use the show wwn switch command to obtain the peer WWN.
	<code>switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00</code>	Disables special frames (default).
Step 2	<code>switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code>	Enables special frames and sets the peer WWN as specified by the profile ID (155).
	<code>switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code>	Disables special frames (default).
Step 3	<code>switch(config-if)# no shutdown</code>	Enables the interface.

Configuring Active Connection

Use the **passive-mode** option to configure the required mode for initiating an IP connection. By default, active mode is enabled to actively attempt an IP connection.

If you enable the passive mode, the switch does not initiate a TCP connection and merely waits for the peer to connect to it.



Note

Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection will not be initiated.

To enable the passive mode, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if) # passive-mode</code>	Enable passive mode while attempting a TCP connection.
	<code>switch(config-if) # no passive-mode</code>	Reverts to the factory set default of using the active mode while attempting the TCP connection.
Step 2	<code>switch(config-if) # no shutdown</code>	Enables the interface.

Configuring the Number of TCP Connections

Use the **tcp-connection** option to specify the number of TCP connections from a FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure 1 or 2 TCP connections.

For example, the Cisco PA-FC-1G Fibre Channel port adapter which has only 1 (one) TCP connection interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit and you can change the configuration on the switch using the **tcp-connection 1** command. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, the software handles it gracefully and moves on with just one connection.

To specify the TCP connection attempts, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if) # tcp-connection 1</code>	Specifies the number of TCP connections. Two (2) is the default and the maximum number of TCP connection attempts.
	<code>switch(config-if) # no tcp-connection 1</code>	Reverts to the factory set default of two attempts.
Step 2	<code>switch(config-if) # no shutdown</code>	Enables the interface.

Enabling Time Stamps

Use the **time-stamp** option to enable or disable FCIP time stamps on a packet. The **time stamp** option instructs the switch to discard packets that are outside the specified time. By default, the **time-stamp** option is disabled.

The **acceptable-diff** option specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped. By default if a packet arrives within a 1000 millisecond interval (+ or -1000 milliseconds), that packet is accepted.



Note

If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the “[NTP Configuration](#)” section on page 4-18).

To enable or disable the **time-stamp** option, follow these steps:

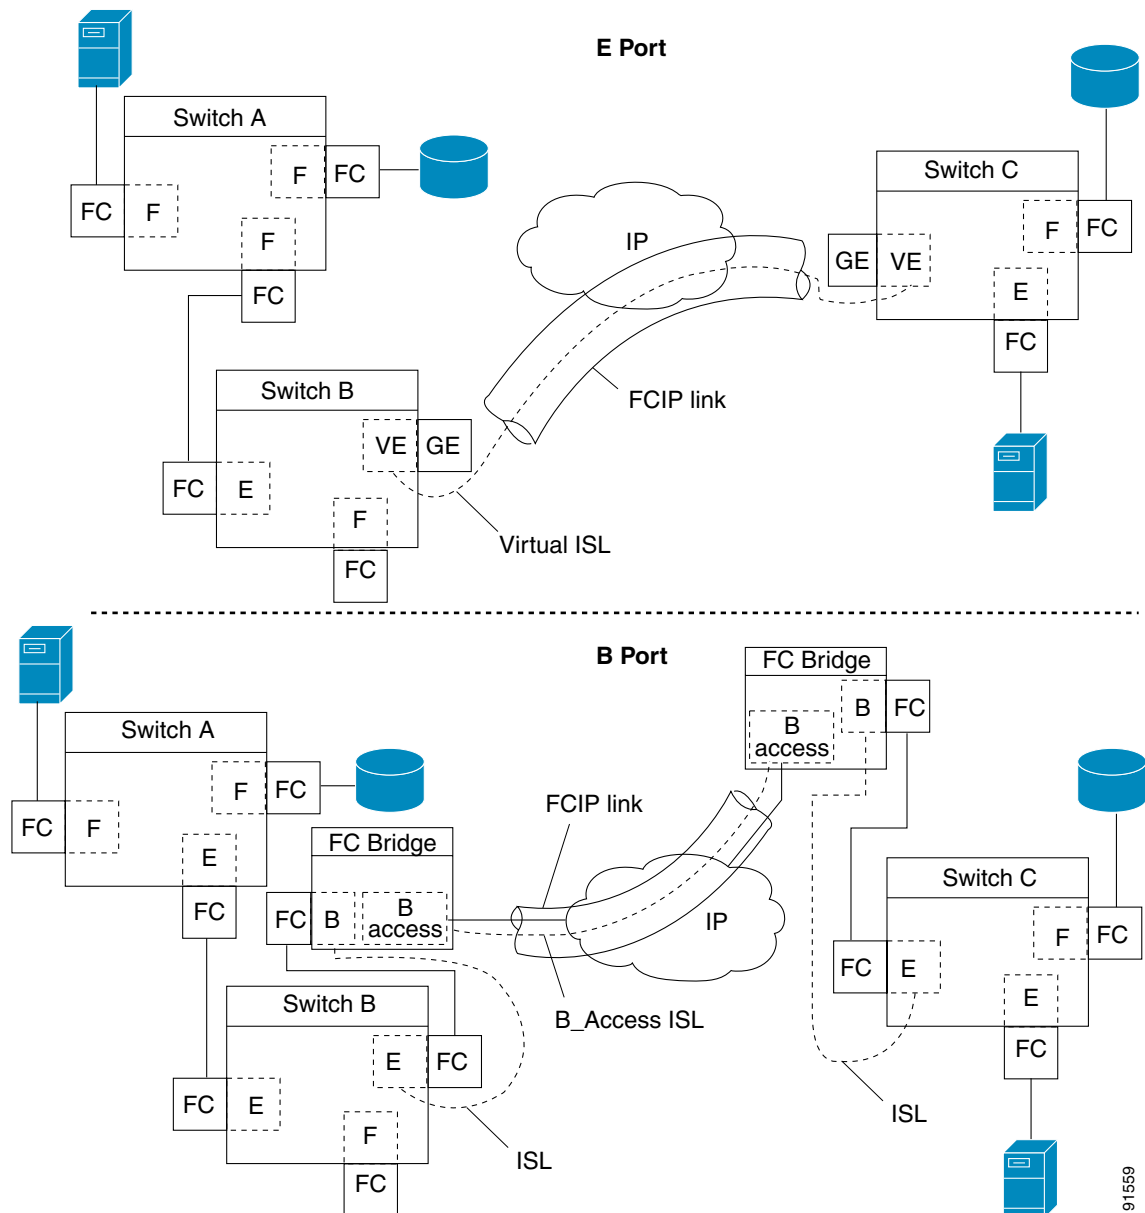
	Command	Purpose
Step 1	<code>switch(config-if) # time-stamp</code> Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link	Enables time stamp checking for received packets with a default acceptable time difference of 1000 milliseconds.
	<code>switch(config-if) # no time-stamp</code>	Disables (default) time stamps.

	Command	Purpose
Step 2	<code>switch(config-if)# time-stamp acceptable-diff 4000</code>	Configures the acceptable time within which a packet is accepted. The default difference is a 1000 millisecond interval from the network time. The valid range is from 500 to 10,000 milliseconds.
	<code>switch(config-if)# no time-stamp acceptable-diff 500</code>	Deletes the configured time difference and reverts the difference to factory defaults.
Step 3	<code>switch(config-if)# no shutdown</code>	Enables the interface.

B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. [Figure 22-11](#) depicts a typical SAN extension over an IP network.

Figure 22-11 FCIP B Port and Fibre Channel E Port



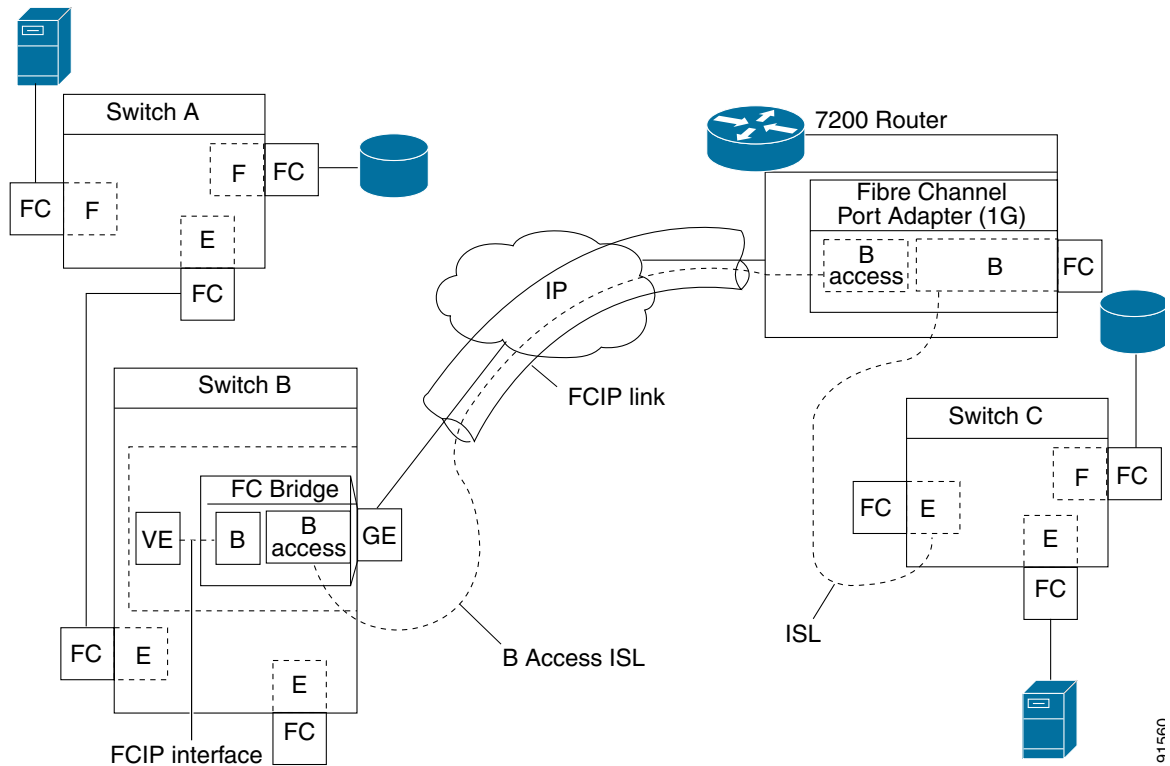
B ports bridge Fibre Channel traffic from one E port to a remote E port without participating in fabric-related activities such as principal switch election, Domain ID assignment, and Fibre Channel routing (FSPF). For example, Class F traffic entering a SAN extender does not interact with the B port.

The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information which ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over a FCIP link, B ports use a B access ISL.*

The IPS module supports FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding virtual B port connects to an virtual E port which completes the end-to-end E port connectivity requirement (see [Figure 22-12](#)).

Figure 22-12 FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

Configuring B Ports

When a FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

To enable B port mode, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# bport</code>	Enables B port mode on the FCIP interface.
	<code>switch(config-if)# no bport</code>	Reverts to E port mode on the FCIP interface (default).
Step 2	<code>switch(config-if)# bport-keepalive</code>	Enables the reception of keepalive responses sent by a remote peer.
Step 3	<code>switch(config-if)# no bport-keepalive</code>	Disables the reception of keepalive responses sent by a remote peer (default).

E Port Configurations

All configuration commands that apply to E ports, also apply to FCIP interfaces. The following features are also available FCIP interfaces:

- VSANs (see [Chapter 9, “Configuring and Managing VSANs”](#))
 - FCIP interfaces can be a member of any VSAN.
- Trunk mode (see [Chapter 11, “Configuring Trunking”](#))
 - Trunk mode can be configured.
 - Trunk allowed VSANs can be configured
- PortChannels (see [Chapter 12, “Configuring PortChannels”](#))
 - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
 - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
- FSPF (see [Chapter 19, “Configuring Fibre Channel Routing Services and Protocols”](#))
- Fibre Channel domains (fcdomains—see [Chapter 24, “Configuring Domain Parameters”](#))
- Zone merge (see [Chapter 13, “Configuring and Managing Zones”](#))
 - Importing the zone database from the adjacent switch.
 - Exporting the zone database from the adjacent switch.

Configuring FCIP Write Acceleration

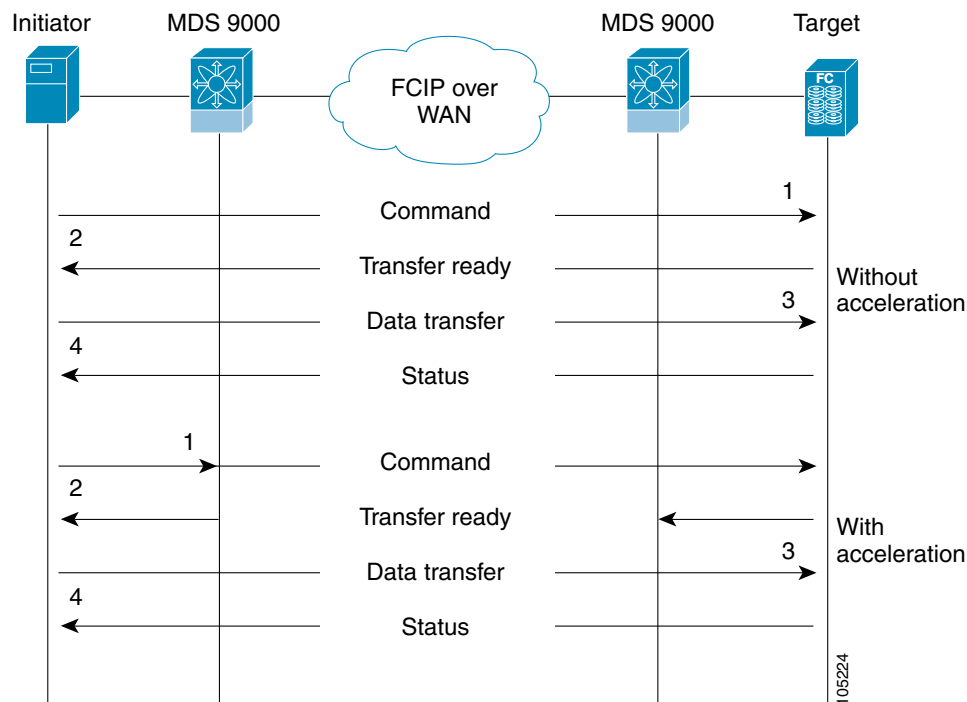
The FCIP Write Acceleration feature in SAN-OS 1.3(3) enables you to significantly improve application performance when storage traffic is routed over wide area networks using FCIP. When FCIP Write Acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for the command to transfer ready acknowledgement (see [Figure 22-13](#)).



Note

The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tunnel will not initialize.

Figure 22-13 FCIP Link Write Acceleration



In [Figure 22-13](#), some data sent by the host is queued on the target before the target issues a Transfer Ready. This way the actual write operation may be done in a less time than the write operation without the write acceleration feature being enabled.



Tip

FCIP write acceleration will not work if the FCIP port is part of a PortChannel or if there are multiple paths with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.

To enable write acceleration, follow these steps:

	Command	Purpose
Step 1	switch1# config terminal switch(config)#	Enters configuration mode.
Step 2	switch1(config)# interface fcip 51 switch1(config-if)#	Creates a FCIP interface (51).
Step 3	switch1(config-if)# write-accelerator	Enables write acceleration.
	switch1(config-if)# no write-accelerator	Disables write acceleration (default).

Enabling FCIP Compression

The FCIP compression feature introduced in Release 1.3(x) allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled.

This feature uses the Lempel-Zif-Stac (LZS) compression algorithm to compress packets.

The **high-throughput** mode allows faster compression but the compression ratio may be lower. The **high-comp-ratio** mode allows a higher compression ratio, but the throughput may be lower.

To enable FCIP compression, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface fcip 51 switch(config-if)#	Creates a FCIP interface (51).
Step 3	switch(config-if)# ip-compression high-throughput	Enables faster compression.
	switch(config-if)# ip-compression high-comp-ratio	Enables a better compression ratio.
	switch(config-if)# no ip-compression	Disables (default) the FCIP compression feature.

Displaying FCIP Information

Use the **show interface** commands to view the summary, counter, description, and status of the FCIP link. Use the output of these commands to verify the administration mode, the interface status, the operational mode, the related VSAN ID, and the profile used. See Examples 22-14 to 22-19.

Example 22-14 Displays the FCIP Interface

```
switch# show interface fcip 3
fcip3 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:ca:00:05:30:00:07:1e
  Peer port WWN is 20:ca:00:00:53:00:18:1e
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  vsan is 1
  Trunk vsans (allowed active) (1,10)
  Trunk vsans (operational) (1)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) (10)
  Trunk vsans (initializing) ()
```



```

Using Profile id 3 (interface GigabitEthernet4/3)
Peer Information
  Peer Internet address is 43.1.1.1 and port is 3225
  Special Frame is disabled
Maximum number of TCP connections is 2
Time Stamp is disabled
B-port mode disabled
TCP Connection Information
  2 Active TCP connections
    Control connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65532
    Data connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65534
  30 Attempts for active connections, 0 close of connections
TCP Parameters
  Path MTU 1500 bytes
  Current retransmission timeout is 300 ms
  Round trip time: Smoothed 10 ms, Variance: 5
  Advertised window: Current: 122 KB, Maximum: 122 KB, Scale: 1
  Peer receive window: Current: 114 KB, Maximum: 114 KB, Scale: 1
  Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
808 frames input, 75268 bytes
  808 Class F frames input, 75268 bytes
  0 Class 2/3 frames input, 0 bytes
  0 Error frames timestamp error 0
806 frames output, 74712 bytes
  806 Class F frames output, 74712 bytes
  0 Class 2/3 frames output, 0 bytes
  0 Error frames 0 reass frames

```

Example 22-15 Displays Detailed FCIP Interface Counter Information

```

switch# show interface fcip 3 counters
fcip3
TCP Connection Information
  2 Active TCP connections
    Control connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65532
    Data connection: Local 43.1.1.2:3225, Remote 43.1.1.1:65534
  30 Attempts for active connections, 0 close of connections
TCP Parameters
  Path MTU 1500 bytes
  Current retransmission timeout is 300 ms
  Round trip time: Smoothed 10 ms, Variance: 5
  Advertised window: Current: 122 KB, Maximum: 122 KB, Scale: 1
  Peer receive window: Current: 114 KB, Maximum: 114 KB, Scale: 1
  Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
5 minutes input rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
5 minutes output rate 64 bits/sec, 8 bytes/sec, 0 frames/sec
814 frames input, 75820 bytes
  814 Class F frames input, 75820 bytes
  0 Class 2/3 frames input, 0 bytes
  0 Error frames timestamp error 0
812 frames output, 75264 bytes
  812 Class F frames output, 75264 bytes
  0 Class 2/3 frames output, 0 bytes
  0 Error frames 0 reass frames

```

Example 22-16 Displays Brief FCIP Interface Counter Information

```

switch# show interface fcip 3 counters brief
-----
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)

```

	Rate	Total	Rate	Total
	Mbits/s	Frames	Mbits/s	Frames
fcip3	9	0	9	0

Example 22-17 Displays the FCIP Interface Description

```
switch# show interface fcip 51 description
FCIP51
    Sample FCIP interface
```

Example 22-18 Displays FCIP Profiles

```
switch# show fcip profile
```

ProfileId	Ipaddr	TcpPort
1	10.10.100.150	3225
2	10.10.100.150	3226
40	40.1.1.2	3225
100	100.1.1.2	3225
200	200.1.1.2	3225

Example 22-19 Displays the Specified FCIP Profile Information

```
switch# show fcip profile 7
FCIP Profile 7
    Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
    Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discovery is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Send buffer size is 0 KB
        Maximum allowed bandwidth is 1000000 kbps
        Minimum available bandwidth is 15000 kbps
        Estimated round trip time is 1000 usec
```

[Example 22-20](#) provides a sample output of FCIP counters when the write-acceleration feature is enabled.

Example 22-20 Displays IP Compression Counters in a FCIP Interface

```
switch# show interface fcip 8 counters
fcip8
    TCP Connection Information
        2 Active TCP connections
            Control connection: Local 10.2.2.1:3225, Remote 10.2.2.2:65439
            Data connection: Local 10.2.2.1:3225, Remote 10.2.2.2:65441
        4 Attempts for active connections, 0 close of connections
    TCP Parameters
        Path MTU 1500 bytes
        Current retransmission timeout is 200 ms
        Round trip time: Smoothed 2 ms, Variance: 1
        Advertized window: Current: 14 KB, Maximum: 14 KB, Scale: 9
        Peer receive window: Current: 14 KB, Maximum: 17 KB, Scale: 9
        Congestion window: Current: 10 KB, Slow start threshold: 112 KB
```

```

5 minutes input rate 760 bits/sec, 95 bytes/sec, 0 frames/sec
5 minutes output rate 912 bits/sec, 114 bytes/sec, 0 frames/sec
  9379771 frames input, 16906568212 bytes
    638 Class F frames input, 58752 bytes
    9379133 Class 2/3 frames input, 16906509460 bytes
    7908669 Reass frames
    0 Error frames timestamp error 0
  9229787 frames output, 16569073984 bytes
    638 Class F frames output, 60128 bytes
    9229149 Class 2/3 frames output, 16569013856 bytes
    0 Error frames
Write Accelerator statistics
  18609558 packets in      10219163 packets out
  0 frames dropped 0 CRC errors
  0 rejected due to table full
  0 ABTS sent      6 ABTS received
  0 tunnel synchronization errors
  485136 writes recd      485136 XFER_RDY sent (host)
  485136 XFER_RDY rcvd (host)
  0 XFER_RDY not proxied due to flow control (host)
  0 bytes queued for sending
  0 estimated bytes queued on the other side for sending
  0 times TCP flow ctrl (target)
  0 bytes current TCP flow ctrl (target)
IP compression statistics
  10044 rxbytes      0 rxbytes compressed
  10044 txbytes      6460 txbytes compressed

```

FCIP High Availability

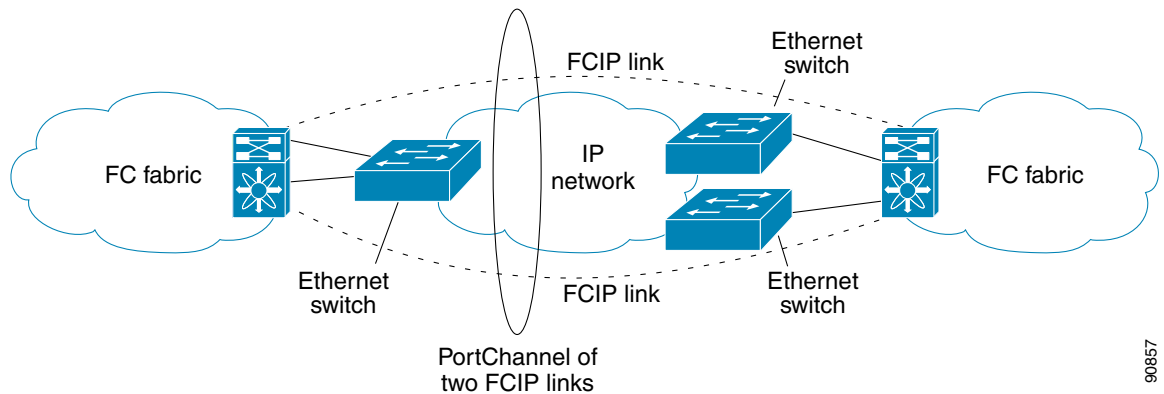
The following high availability solutions are available for FCIP configurations:

- [Fibre Channel PortChannels, page 22-40](#)
- [FSPF, page 22-40](#)
- [VRRP, page 22-41](#)
- [Ethernet PortChannels, page 22-41](#)

Fibre Channel PortChannels

Figure 22-14 provides an example of a PortChannel-based load balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 22-14 PortChannel Based Load Balancing



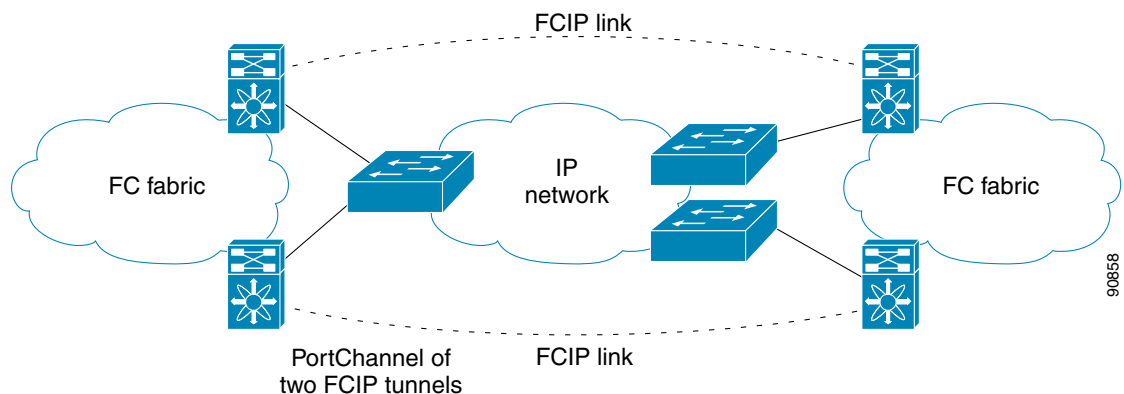
The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the PortChannel should be across the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the PortChannel.

FSPF

Figure 22-15 displays a FSPF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 22-15 FSPF-Based Load Balancing



The following characteristics set FSPF solutions apart from other solutions:

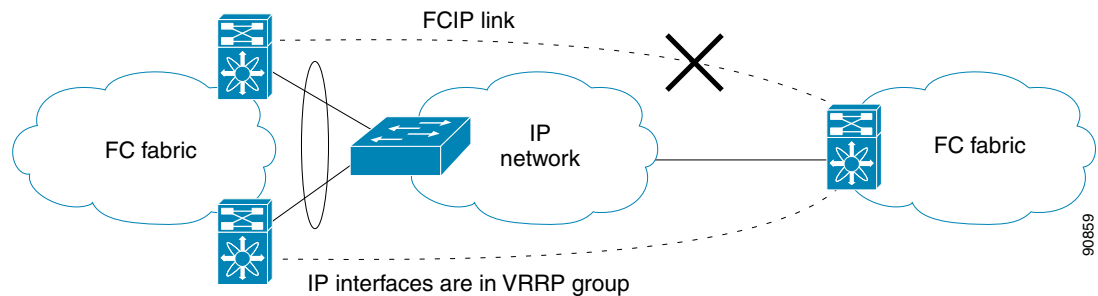
- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.

- The Fibre Channel traffic is load balanced across the FCIP link.

VRRP

Figure 22-16 displays a VRRP-based high availability FCIP configuration example. This configuration, requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

Figure 22-16 VRRP-Based High Availability



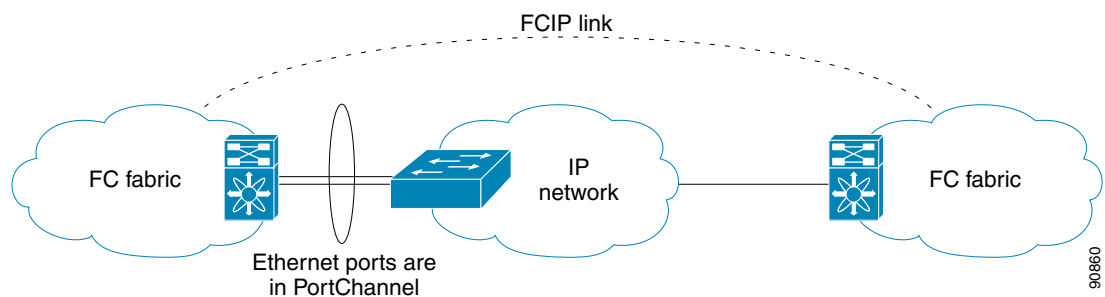
The following characteristics set VRRP solutions apart from other solutions:

- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

Ethernet PortChannels

Figure 22-17 displays a Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 22-17 Ethernet PortChannel-Based High Availability



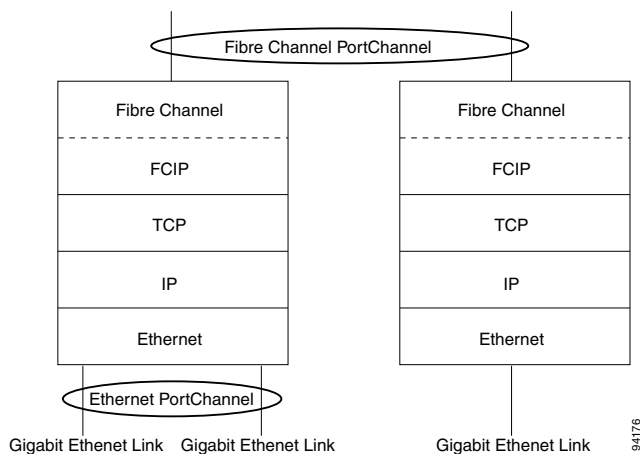
The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- The Gigabit Ethernet link level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appears like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.

Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer Ethernet-level redundancy, Fibre Channel PortChannels offer (E)ISL-level redundancy. FCIP is unaware of any Ethernet PortChannels or Fibre Channel PortChannels. Fibre Channel PortChannels are unaware of any Ethernet PortChannels, and there is no mapping between the two (see , [page 22-42](#)).

Figure 22-18 PortChannels at the Fibre Channel and Ethernet Levels



To configure Fibre Channel PortChannels, see [Chapter 12, “Configuring PortChannels.”](#) To configure Ethernet PortChannels, refer to the [“Configuring Ethernet PortChannels”](#) section on [page 22-14](#).

Configuring iSCSI

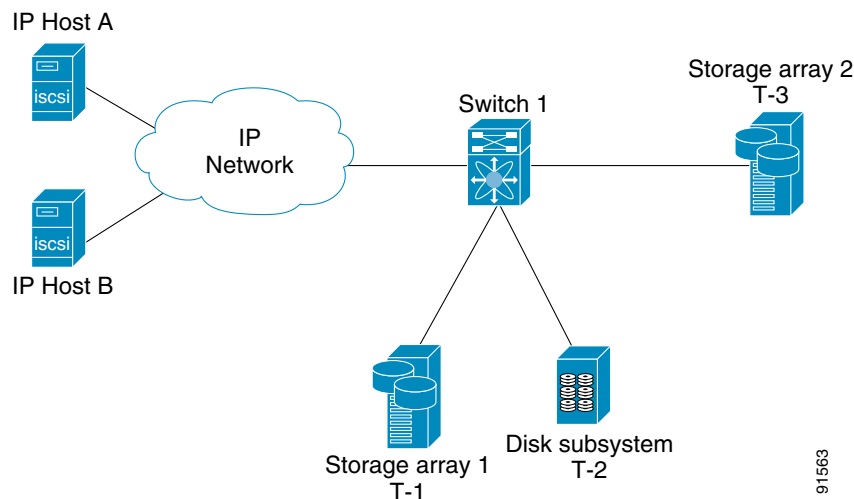
This section includes the following topics:

- [About iSCSI, page 22-43](#)
- [Enabling iSCSI, page 22-45](#)
- [Presenting Fibre Channel Targets as iSCSI Targets, page 22-46](#)
- [iSCSI Virtual Target Configuration Examples, page 22-50](#)
- [Presenting iSCSI Hosts as Virtual Fibre Channel Hosts, page 22-52](#)
- [Configuring iSCSI Proxy Initiators, page 22-56](#)
- [Access Control in iSCSI, page 22-58](#)
- [iSCSI User Authentication, page 22-60](#)
- [Assigning VSAN Membership to iSCSI Hosts, page 22-55](#)
- [Advanced iSCSI Configuration, page 22-62](#)
- [Displaying iSCSI Information, page 22-64](#)
- [iSCSI High Availability, page 22-75](#)
- [iSCSI Authentication Setup Guidelines, page 22-78](#)

About iSCSI

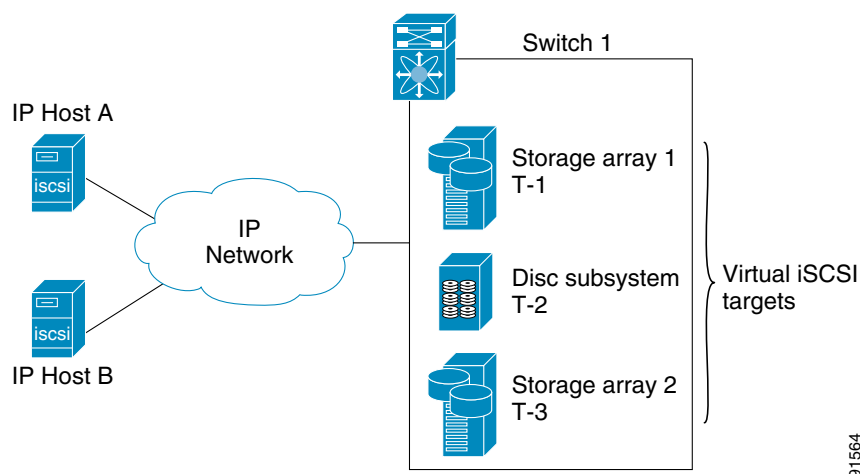
The IPS module provides transparent SCSI routing by default. IP hosts using iSCSI protocol can transparently access targets on the Fibre Channel network. [Figure 22-19](#) provides an example of a typical configuration of iSCSI hosts with access to a Fibre Channel SAN.

Figure 22-19 Typical IP to Fibre Channel SAN Configuration



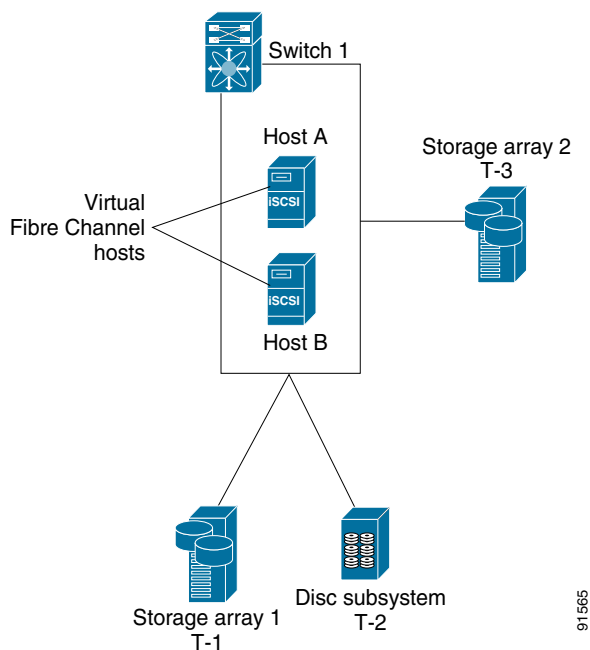
The IPS module enables you to create virtual iSCSI targets and maps them to physical Fibre Channel targets available in the Fibre Channel SAN. It presents the Fibre Channel targets to IP hosts as if the physical targets were attached to the IP network (see Figure 22-20).

Figure 22-20 iSCSI View



In conjunction with presenting Fibre Channel targets to iSCSI hosts, the IPS module presents each iSCSI host as a Fibre Channel host (in transparent mode), i.e. Host Bus Adaptor (HBA) to the Fibre Channel storage device. The storage device responds to each IP host as if it were a Fibre Channel host connected to the Fibre Channel network (see Figure 22-21).

Figure 22-21 Fibre Channel SAN View



Note

Refer to the IETF standards for IP storage at <http://www.ietf.org>, for information on the iSCSI protocol.

Enabling iSCSI

To begin configuring the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the iSCSI feature are only available when iSCSI is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

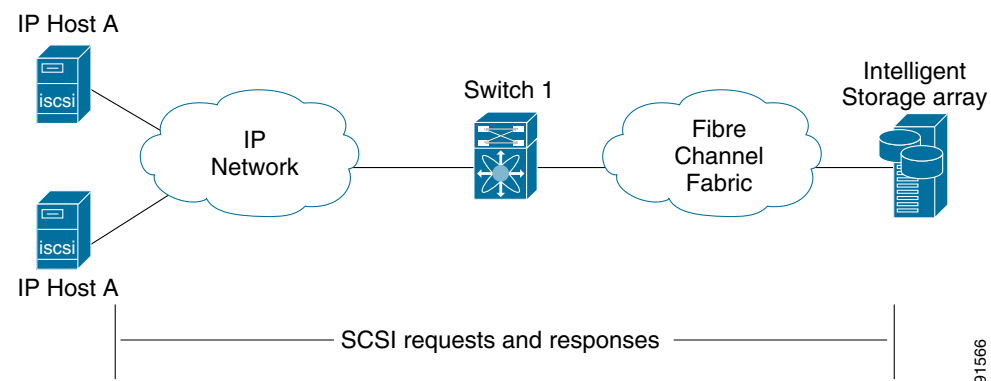
To enable iSCSI on any participating switch, follow these steps:

	Command	Purpose
Step 1	switch# conf t	Enters configuration mode.
Step 2	switch(config)# iscsi enable	Enables iSCSI on that switch.
	switch(config)# no iscsi enable	Disables (default) iSCSI on that switch.

Routing iSCSI Requests and Responses

The iSCSI feature consists of routing iSCSI requests and responses between hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see Figure 22-22).

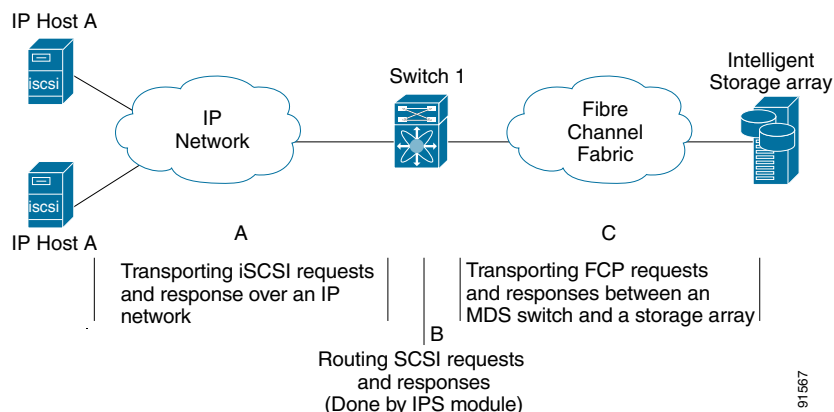
Figure 22-22 Routing iSCSI Requests and Responses for Transparent iSCSI Routing



Each iSCSI host that requires access to storage via the IPS module needs to have a compatible iSCSI driver installed. (The CCO website at <http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml> provides a list of compatible drivers). Using iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver for a peripheral channel in the host. From the storage device perspective, each IP host appears as a Fibre Channel host.

Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions (see Figure 22-22):

- Transporting iSCSI requests and responses over an IP network between hosts and the IPS module.
- Routing SCSI requests and responses between hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). This routing is performed by the IPS module.
- Transporting FCP requests or responses between the IPS module and Fibre Channel storage devices.

Figure 22-23 Transparent SCSI Routing Actions**Note**

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN.

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module presents physical Fibre Channel targets as iSCSI targets allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- **Dynamic Importing**—used if all logical units (LUs) in all Fibre Channel storage targets are made available to iSCSI hosts (subject to VSAN and zoning).
- **Static Importing**—used if iSCSI hosts are restricted to subsets of LUs in the Fibre Channel targets and additional iSCSI access control is needed (see the [“Access Control in iSCSI”](#) section on page 22-58). Also, static import allows automatic failover if the Fibre Channel targets’ LU is reached by redundant Fibre Channel ports (see the [“High Availability Static Target Importing”](#) section on page 22-49).

**Note**

The IPS module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have a configured name. Targets that are not statically imported are advertised with the name created by the conventions explained in this section.

Dynamic Importing

To enable dynamic importing of Fibre Channel targets into iSCSI, use the **iscsi import target fc** command.

The IPS module maps each physical Fibre Channel target port as one iSCSI target. That is, all LU accessible via the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the storage target.

For example, if an iSCSI target was created for Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0 through 2, those LUNs would become available to an IP host as LUNs 0 through 2 as well.

**Note**

If you have configured a switch name, then the switch name will be used instead of the management IP address. If you have not configured a switch name, the management IP address will be used.

The iSCSI target node name is created automatically using the iSCSI qualified name (IQN) format. The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.

The IPS module creates an IQN formatted iSCSI node name using the following conventions:

- IPS ports that are not part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
```

- IPS ports that are part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
```

- Ports that are part of a PortChannel use this format:

```
iqn.1987-02.com.cisco:05.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>
```

**Note**

With this format, each IPS port in a Cisco MDS 9000 Family switch creates a different iSCSI target node name for the same Fibre Channel target.

To dynamically import Fibre Channel targets, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi import target fc	IPS modules dynamically import each Fibre Channel target in the Fibre Channel SAN to the IP network. The automatically-created iSCSI target node names use the IQN format. Note Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms.

Static Importing

You can manually (statically) create an iSCSI target and assign a node name to it. A statically-mapped iSCSI target can either contain the whole FC target port, or it can contain one or more LUs from a Fibre Channel target port.

To create a static iSCSI virtual target for the entire Fibre Channel target port, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator switch(config- (iscsi-tgt))#	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.

Step 3

Command	Purpose
<pre>switch(config-(iscsi-tgt))# pwwn 26:00:01:02:03:04:05:06</pre>	<p>Maps a virtual target node to a Fibre Channel target. One iSCSI target cannot contain more than one Fibre Channel target.</p> <p>Don't specify the LUN if you wish to map the whole Fibre Channel target to an iSCSI target. All Fibre Channel target LUNs are exposed to iSCSI.</p> <p>Use the LUN option to map different Fibre Channel LUNs to different iSCSI virtual targets. If you have already mapped the whole Fibre Channel target, you will not be able to use this option.</p>
<pre>switch(config-(iscsi-tgt))# pwwn 26:00:00:00:00:11:00:11 fc-lun 1 iscsi-lun 1</pre>	Maps a virtual target using LUN mapping options.

Refer to [“iSCSI-Based Access Control” section on page 22-59](#) for more information on controlling access to statically-imported targets.

For multiple interfaces configured with iSNS (see the [“Configuring Storage Name Services” section on page 22-91](#)), a different static virtual target name has to be created for each interface tagged to an iSNS profile and each static virtual target must be advertised only from one interface (see the [“Advertising iSCSI Targets” section on page 22-49](#))

Advertising iSCSI Targets

You can limit the Gigabit Ethernet interfaces over which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

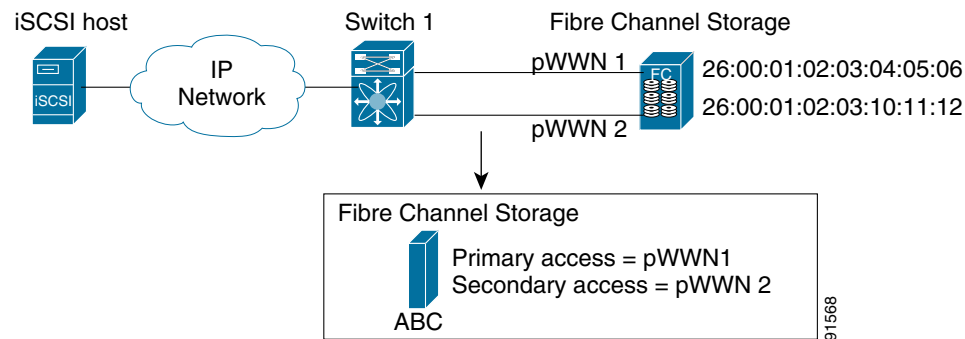
To configure a specific interface that should be advertised but a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-(iscsi-tgt))# advertise interface GigabitEthernet 2/5</code>	Advertises the virtual target only on the specified interface. By default, it is advertised on all interfaces in all IPS modules.
	<code>switch(config-(iscsi-tgt))# no advertise interface GigabitEthernet 2/5</code>	Removes this interfaces from the list of interfaces from which this target is advertised.

High Availability Static Target Importing

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port. iSCSI transparently switches to using the secondary port without impacting the iSCSI host. All other I/O are terminated with check condition status and the host retries the IO. If both the primary and secondary pWWNs are available, then both pWWNs can be used—each session may use either pWWN (see [Figure 22-24](#)).

Figure 22-24 Static Target Importing through Two Fibre Channel Ports



In [Figure 22-24](#), you can create a virtual iSCSI target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

To create a static iSCSI virtual target, follow these steps:

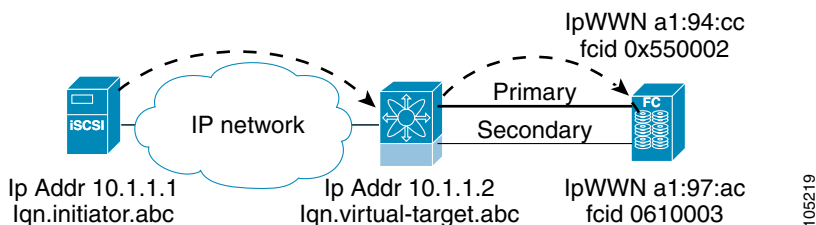
	Command	Purpose
Step 1	<code>switch# config terminal</code> <code>switch(config)#</code>	Enters configuration mode.
Step 2	<code>switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator</code>	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.
Step 3	<code>switch(config-(iscsi-tgt))# pwwn 26:00:01:02:03:04:05:06</code> <code>secondary-pwwn 26:00:01:02:03:10:11:12</code>	Configures the primary and secondary ports for this virtual target.

Configuring the Trespass Feature

In addition to the high availability of statically imported iSCSI targets, the trespass feature is available (effective Release 1.3(x)) to enable the export of LUs, on an active port failure, from the active to the passive port of a statically imported iSCSI target.

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N-ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the **trespass** command be issued, to export the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the MDS issues a trespass command to the target to export the LUs on the new active port. The iSCSI session switches to use the new active port and the exported LUs are accessed over the new active port (see [Figure 22-25](#)).

Figure 22-25 Virtual Target with an Active Primary Port



To enable the Trespass feature for a static iSCSI virtual target, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name 1987-02.com.cisco.initiator switch(config- (iscsi-tgt))#	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.
Step 3	switch(config- (iscsi-tgt))# pwwn 50:00:00:a1:94:cc secondary-pwwn 50:00:00:a1:97:ac	Maps a virtual target node to a Fibre Channel target and configures a secondary pWWN
Step 4	switch(config- (iscsi-tgt))# trespass	Enables the trespass feature.
	switch(config- (iscsi-tgt))# no trespass	Disables the trespass feature (default).

To verify the trespass issue the **show iscsi virtual-target** command in EXEC mode:

```
switch# show iscsi virtual-target iqn.1987-02.com.cisco.initiator
target: 1987-02.com.cisco.initiator
  Port WWN 00:00:00:00:00:00:00:00
  Configured node
  all initiator permit is disabled
  trespass support is enabled
```

iSCSI Virtual Target Configuration Examples

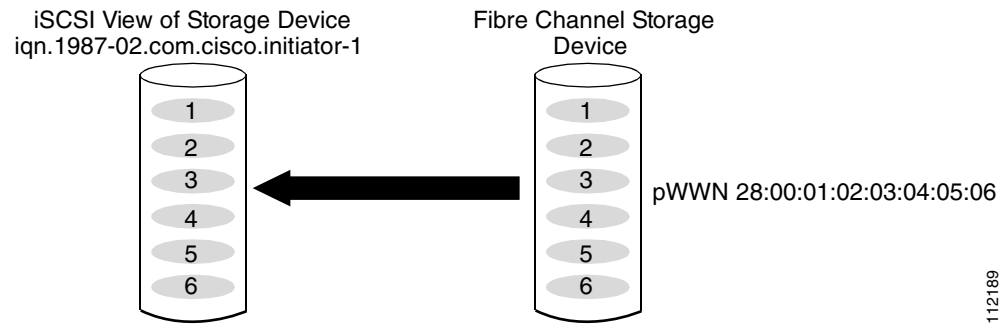
This section provides three examples of virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as a virtual iSCSI target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see [Figure 22-26](#)).

```
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-1
pWWN 28:00:01:02:03:04:05:06
```

Figure 22-26 Assigning iSCSI Node Names



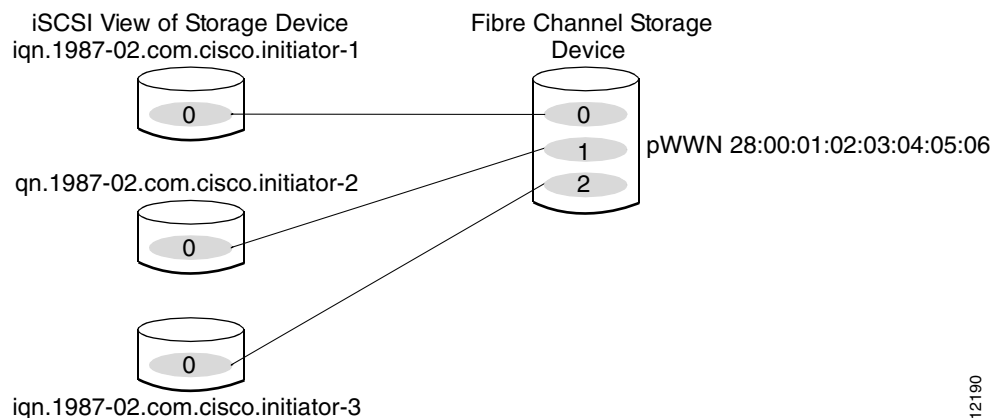
112189

Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN (see [Figure 22-27](#)).

```
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-1
pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-2
pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-3
pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
```

Figure 22-27 Mapping LUNs to iSCSI a Node Name



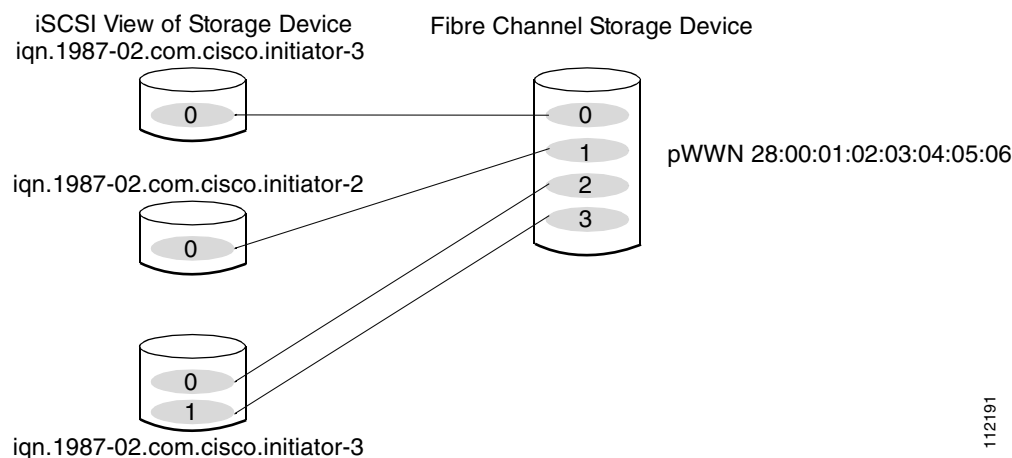
112190

Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see [Figure 22-28](#)).

```
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-1
  pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-2
  pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.initiator-3
  pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
  pWWN 28:00:01:02:03:04:05:06 fc-lun 3 iscsi-lun 1
```

Figure 22-28 Mapping LUNs to Multiple iSCSI Node Names



112/191

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The iSCSI hosts are mapped to virtual Fibre Channel hosts in one of two ways (see [Figure 22-21](#)):

- **Dynamic Mapping** (default)—used if no access control is done on the Fibre Channel target. An iSCSI host may use different pWWNs each time it connects to a Fibre Channel target.
- **Static Mapping**—used if an iSCSI host should always have the same pWWN or nWWN each time it connects to a Fibre Channel target.

Dynamic Mapping

When an iSCSI host connects to the IPS module using the iSCSI protocol, a virtual N port is created for the host. The nWWNs and pWWNs are dynamically allocated from the switch's Fibre Channel WWN pool. The IPS module registers this N port in the Fibre Channel SAN. The IPS module continues using that nWWN and pWWN to represent this iSCSI host until it no longer has a connection to any iSCSI target via that IP storage port.

At that point, the virtual Fibre Channel host is taken offline from the Fibre Channel SAN and the nWWNs and pWWNs are released back to the switch's Fibre Channel WWN pool. These addresses become available for assignment to other iSCSI hosts requiring access to Fibre Channel SANs.

When a dynamically mapped iSCSI initiator has multiple sessions to multiple Fibre Channel targets, each session can use the same pWWN and nWWN as long as it uses the same node name in the iSCSI login message. If the host has multiple network interfaces (each having different IP addresses), and you want each IP address to be treated as different iSCSI initiator hosts, then use the **switchport initiator id ip-address** command to identify each IP address as an iSCSI initiator.

Identifying Initiators

An iSCSI initiator is identified in one of two ways:

- By iSCSI node name (**switchport initiator id name** command)—an initiator with multiple IP addresses (multiple interface cards—NICs or multiple network interfaces) has one virtual N port, assuming it uses the same iSCSI initiator name to iSCSI targets from all interfaces.
- By IP address (**switchport initiator id ip-address** command)—a virtual N port is created for each IP address it uses to login to iSCSI targets.

By default, the switch uses the iSCSI node name to identify the initiator.

To identify the initiator using the IP address, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface iscsi 4/1 switch(config-if)#	Selects the iSCSI interface on the switch that identifies all the initiators.
Step 3	switch(config-if)# switchport initiator id ip-address	Identifies the iSCSI initiator based on the IP address.
	switch(config-if)# switchport initiator id name	Identifies the iSCSI initiator based on the initiator node name.

Static Mapping

With dynamic mapping, each time the iSCSI host connects to the IPS module a new Fibre Channel N port is created and the nWWNs and pWWNs allocated for this N port may be different. Use the static mapping method to obtain the same nWWN and pWWNs for the iSCSI host each time it connects to the IPS module.

You can implement static mapping in one of two ways: system assignment or manual assignment.

- System assignment—When a static mapping configuration is created, one nWWN and/or one or more pWWNs are allocated from the switch's Fibre Channel WWN pool and the mapping is kept permanent. This assignment uses the **system-assign** option.
- Manual assignment—You can specify your own unique WWN by providing them during the configuration process.



Tip

We recommend using the **system-assign** option. If you manually assign a WWN, you must ensure its uniqueness (see the [“Configuring World Wide Names”](#) section on page 29-18).

Static mapping can be used on the IPS module to access intelligent Fibre Channel storage arrays that have access control and LUN mapping/masking configuration based on the initiator's pWWNs and/or nWWNs.

**Note**

If an iSCSI host connects to multiple IPS ports, each port independently creates one virtual N port for the host. If static mapping is used, enough pWWNs should be configured for as many IPS ports to which a host connects.

To configure static mapping (using the **name** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config- (iscsi-init))#	Configures an iSCSI initiator using the iSCSI name of the initiator node. The maximum name length is restricted to 223 alphanumeric characters. The minimum length is 16.
	switch(config)# no iscsi initiator name iqn.1987-02.com.cisco.initiator	Deletes the configured iSCSI initiator.

To configure static mapping (using the **ip-address** option) for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator ip address 10.50.0.0 switch(config- (iscsi-init))#	Configures an iSCSI initiator. using the IP address of the initiator node.
	switch(config)# no iscsi initiator ip address 10.50.0.0	Deletes the configured iSCSI initiator.

To assign the WWN for an iSCSI initiator, follow these steps:

	Command	Purpose
Step 1	switch(config- (iscsi-init))# static nWWN system-assign	Uses the switch's WWN pool to allocate the nWWN for this iSCSI initiator and keeps it persistent.
	switch(config- (iscsi-init))# nWWN 20:00:00:05:30:00:59:11	Assigns the user provided WWN as nWWN for the iSCSI initiator. You can only specify one nWWN for each iSCSI node.
Step 2	switch(config- (iscsi-init))# static pWWN system-assign 2	Uses the switch's WWN pool to allocate two pWWNs for this iSCSI initiator and keeps it persistent. The range is from 1 to 64.
	switch(config- (iscsi-init))# pWWN 21:00:00:20:37:73:3b:20	Assigns the user provided WWN as pWWN for the iSCSI initiator.

**Note**

If a system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is backed up to an ASCII file the system-assigned WWNs are also saved. Subsequently if you issue a **write erase** command, you must manually delete the WWN configuration from the ASCII file.

Making the Dynamic Initiator WWN Mapping Static

After a dynamic initiator has already logged in, you may decide to permanently keep the automatically-assigned nWWN/pWWN mapping, so this initiator uses the same mapping the next time it logs in.

To permanently keep the automatically assigned nWWN/pWWN mapping, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi save-initiator name iqn.1987-02.com.cisco.initiator	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose name is specified.
	switch(config)# iscsi save-initiator ip-address 10.10.100.11	Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose IP address is specified.
	switch(config)# no iscsi save-initiator name iqn.1987-02.com.cisco.initiator	Removes the static nWWN and pWWNs mapping for the iSCSI initiator whose name is specified.

Assigning VSAN Membership to iSCSI Hosts

By default, a host is only in VSAN 1 (default VSAN). You can configure an iSCSI host to be a member of one or more VSANs. The IPS module creates one Fibre Channel virtual N port in each VSAN to which the host belongs.

To assign VSAN membership for iSCSI hosts, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-(iscsi-init))#	Configures an iSCSI initiator.
Step 3	switch(config-(iscsi-init))# vsan 3	Assigns the iSCSI initiator node to a specified VSAN.
		Note You can assign this host to one or more VSANs.
	switch(config-(iscsi-init))# no vsan 5	Removes the iSCSI node from the specified VSAN.



Note

When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

Assigning VSANs to a iSCSI Interface

All dynamic iSCSI initiators are members of VSAN 1. The port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. The default port VSAN of an iSCSI interface is VSAN 1, but can be changed using the **vsan *vsan-number* interface *iscsi slot/port*** command in the VSAN database submode. All dynamic iSCSI initiators are member of the port vsan of the iSCSI interface.

To change the default port VSAN, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi interface vsan-membership	Enables you to configure VSAN membership for iSCSI interfaces.
Step 3	switch(config)# vsan database switch(config-vsan-db)#	Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.
Step 4	switch(config-vsan-db)# vsan 2 interface iscsi 2/1	Assigns the membership of the iscsi 2/1 interface to the specified VSAN (VSAN 2).
	switch(config-vsan-db)# no vsan 2 interface iscsi 2/1	Reverts to using the default VSAN as the port VSAN of the iSCSI interface.



Tip

This is a 1.3(x) feature. If you downgrade to an earlier release, be sure to delete any assigned VSAN and to issue the **no iscsi interface vsan-membership** command before performing the downgrade procedure.

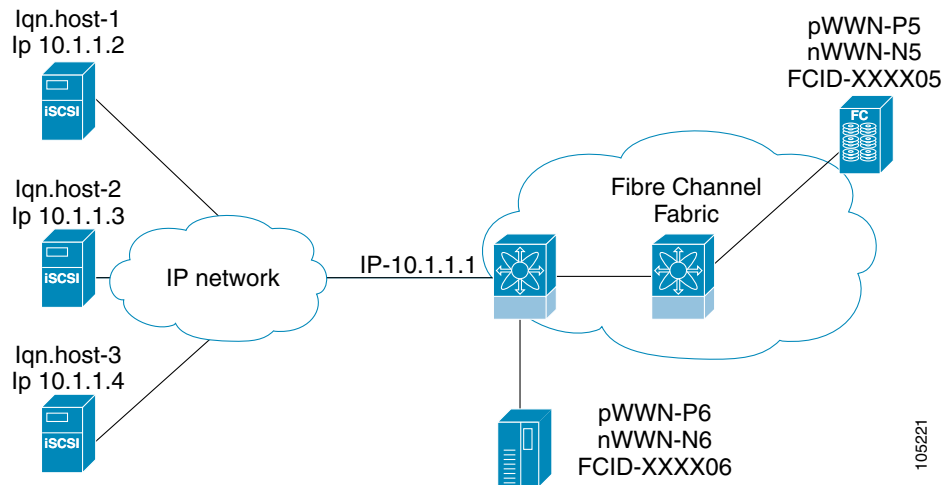
Configuring iSCSI Proxy Initiators



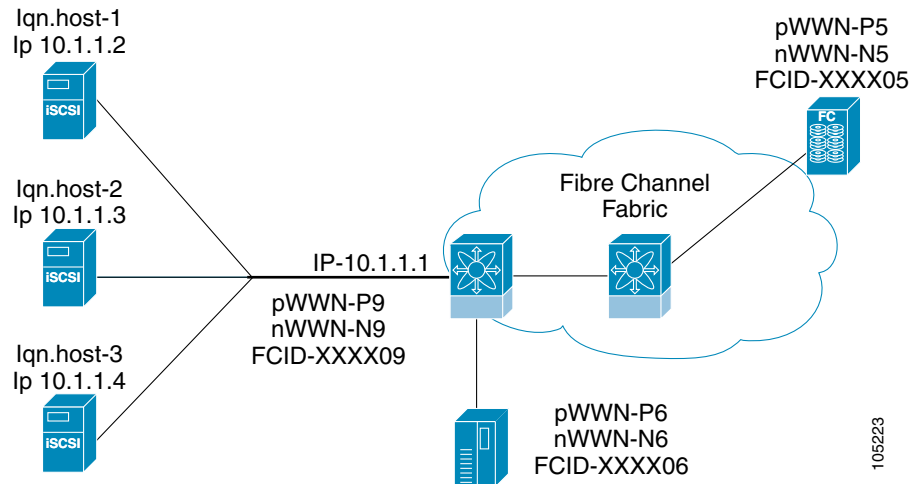
Note

When an interface is in the proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the Fibre Channel interface attributes—the WWN pair and available FCIDs. You cannot configure zoning using iSCSI attributes such as the IP address or the iQN name of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the [“Access Control in iSCSI”](#) section on page 22-58).

By default, each iSCSI initiator appears as one Fibre Channel initiator in transparent mode in the Fibre Channel fabric. For some storage arrays, this appearance requires the initiator's pWWN to be manually configured for access control purposes. This process can be quite cumbersome. The Proxy initiator feature allows all iSCSI initiators to connect through one IPS port making it appear as one Fibre Channel port per VSAN. It simplifies the task of configuring the pWWN for each new initiator on the storage array, and Fibre Channel access control such as zoning. This feature along with static target importing (using LUN mapping) results in the configuration being performed only on the switch when a new iSCSI host is added. On the storage array, all LUNs that will be used by iSCSI initiators are configured to allow access by the proxy initiator's pWWN. From the iSCSI perspective, this configuration is no different from the default mode (see [Figure 22-29](#)).

Figure 22-29 The iSCSI View of a Proxy Initiator

From the Fibre Channel perspective, only one Fibre Channel initiator is visible per VSAN (see [Figure 22-30](#)).

Figure 22-30 The FC View with a Proxy Initiator

To configure the proxy initiator, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface iscsi 4/1 switch(config-if)#	Selects the iSCSI interface on the switch that identifies all the initiators.
Step 3	switch(config-if)# switchport proxy-initiator	Configure the proxy initiator mode using the switch's pWWN/nWWN pool.
	switch(config-if)# no switchport proxy-initiator	Deletes the proxy initiator mode.

	Command	Purpose
Step 4	switch(config-if)# switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22:22	Configures the proxy initiator mode using the specified WWNs.
	switch(config-if)# no switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22:22	Deletes the proxy initiator mode.

To verify the proxy initiator mode configuration, use the **show interface iscsi** command in EXEC mode (see the [“Displaying Proxy Initiator Information”](#) section on page 22-66).

Access Control in iSCSI

You can control access to each statically-mapped iSCSI target by specifying a list of IPS ports on which it will be advertised and specifying a list of iSCSI initiator node names allowed to access it. Fibre Channel zoning-based access control and iSCSI-based access control are the two mechanisms by which access control can be provided for iSCSI. Both methods can be used simultaneously.



Note

This access control is in addition to the existing Fibre Channel access control. The iSCSI initiator has to be in the same VSAN and zone as the physical Fibre Channel target.

Fibre Channel Zoning-Based Access Control

Zoning is an access control mechanism within a VSAN. The switch's zoning implementation extends the VSAN and zoning concepts from the Fibre Channel domain to also cover the iSCSI domain. This extension includes both iSCSI and Fibre Channel features and provides a uniform, flexible access control across a SAN. Static and dynamic are the two Fibre Channel zoning access control mechanisms.

- Static—statically map the iSCSI host to Fibre Channel virtual N port(s). This creates a permanent nWWNs and pWWNs. Next, configure the assigned pWWN into zones, similar to adding a regular Fibre Channel host's pWWN to a zone.
- Dynamic—add the iSCSI host's initiator node name as a member of a zone. When the IP host's Fibre Channel virtual N port is created and the Fibre Channel address (nWWNs and pWWNs) is assigned, Fibre Channel zoning is enforced.

To register an iSCSI initiator in the zone database, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# zone name iSCSIzone vsan 1 switch(config-zone)	Creates a zone name for the iSCSI devices in the IPS module to be included.

	Command	Purpose
Step 3	switch(config-zone)# member symbolic-nodename iqn.1987-02.com.cisco.initiator1	Adds the device as specified by the node name.
	switch(config-zone)# no member iqn.1987-02.com.cisco.initiator1	Deletes the specified device.
	switch(config-zone)# member ip-address 10.50.1.1	Adds the device identified by the IP address.
	switch(config-zone)# no member ip-address 10.50.1.1	Deletes the identified device.
	switch(config-zone)# member pwwn 20:00:00:05:30:00:59:11	Adds the device identified by the port WWN.
	switch(config-zone)# no member pwwn 20:00:00:05:30:00:59:11	Deletes the device identified by the port WWN
	switch(config-zone)# member ip-address 10.50.1.1 255.255.0.0	Adds all devices in the specified IP subnet.
	switch(config-zone)# no member ip-address 10.50.1.1 255.255.0.0	Deletes all devices in the specified IP subnet.

iSCSI-Based Access Control

For static iSCSI targets, you can manually configure a list of iSCSI initiators that are allowed to access it. The iSCSI initiator is identified by the iSCSI node name or the IP address of the iSCSI host.

By default, static virtual iSCSI targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow a virtual iSCSI target to be accessed by all hosts. The initiator access list can contain one or more initiators. Each initiator is identified by one of the following:

- iSCSI node names
- IP addresses
- IP subnets

To configure access control in iSCSI, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator switch(config-(iscsi-tgt))#	Creates the iSCSI target name iqn.1987-02.com.cisco.initiator.
Step 3	switch(config-(iscsi-tgt))# pwwn 26:00:01:02:03:04:05:06 switch(config-(iscsi-tgt))#	Maps a virtual target node to a Fibre Channel target.

Step 4

Command	Purpose
<code>switch(config-(iscsi-tgt))# initiator iqn.1987-02.com.cisco.initiator1 permit</code>	Allows the specified iSCSI initiator node to access this virtual target. You can issue this command multiple times to allow multiple initiators.
<code>switch(config-(iscsi-tgt))# no initiator iqn.1987-02.com.cisco.initiator1 permit</code>	Prevents the specified initiator node from accessing virtual targets.
<code>switch(config-(iscsi-tgt))# initiator ip address 10.50.1.1 permit</code>	Allows the specified IP address to access this virtual target. You can issue this command multiple times to allow multiple initiators.
<code>switch(config-(iscsi-tgt))# no initiator ip address 10.50.1.1 permit</code>	Prevents the specified IP address from accessing virtual targets.
<code>switch(config-(iscsi-tgt))# initiator ip address 10.50.1.1 255.255.255.0 permit</code>	Allows all initiators in this subnetwork to access this virtual target.
<code>switch(config-(iscsi-tgt))# no initiator ip address 10.50.1.1 255.255.255.0 permit</code>	Prevents all initiators in this subnetwork from accessing virtual targets.
<code>switch(config-(iscsi-tgt))# all-initiator-permit</code>	Allows all initiator nodes to access this virtual target.
<code>switch(config-(iscsi-tgt))# no all-initiator-permit</code>	Prevents any initiator from accessing virtual targets (default).

Enforcing Access Control

IPS modules use both iSCSI node name-based and Fibre Channel zoning-based access control lists to enforce access control during iSCSI discovery and iSCSI session creation.

- iSCSI discovery—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section.
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module verifies if the specified iSCSI target (in the session login request) is a static mapped target, and if true, verifies if the IP host's iSCSI node name is allowed to access the target. If the IP host does not have access, its login is rejected.

The IPS module, then creates a Fibre Channel virtual N port (the N port may already exist) for this IP host and does a Fibre Channel name server query for the FCID of the Fibre Channel target pWWN that is being accessed by the IP host. It uses the IP host virtual N port's pWWN as the requester of the name server query. Thus, the name server does a zone-enforced query for the pWWN and responds to the query.

If the FCID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

iSCSI User Authentication

The IPS module supports the iSCSI authentication mechanism to authenticate iSCSI hosts that request access to storage. When iSCSI authentication is enabled, the iSCSI hosts must provide user name and password information each time an iSCSI session is established.

**Note**

Only the Challenge Handshake Authentication Protocol (CHAP) authentication method is supported.

The **aaa authentication iscsi** default command enables aaa authentication for the iSCSI host. If no authentication is configured, local authentication is used. You can use RADIUS authentication (see the “Configuring RADIUS” section on page 16-6) or TACACS+ authentication (see the “Configuring TACACS+” section on page 16-10).

To configure AAA authentication for an iSCSI user, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# aaa authentication iscsi default group RadServerGrp	Uses RADIUS servers which are added in the group called RadServerGrp for the iSCSI authentication method.
	switch(config)# aaa authentication iscsi default group TacServerGrp	Uses TACACS+ servers which are added in the group called TacServerGrp for the iSCSI authentication method.
	switch(config)# aaa authentication iscsi default local	Uses the local password database for iSCSI CHAP authentication.
	switch(config)# iscsi authentication none	Specifies no authentication configuration
	switch(config)# iscsi authentication chap-none	Specifies that both CHAP or no authentication is allowed. Use this option to override the global configuration which might have been configured to allow only one option—either CHAP or none—not both.

Authentication Mechanism

During an iSCSI login, both the iSCSI initiator and target have the option to authenticate each other. By default, the IPS module allows either CHAP authentication or no authentication from iSCSI hosts. If CHAP authentication should always be used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used, issue the **iscsi authentication none** command.

**Note**

The authentication for a Gigabit Ethernet interface or subinterface configuration overrides the authentication for the global interface configuration.

To configure the authentication mechanism for iSCSI, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# iscsi authentication chap	Configures CHAP as the default authentication mechanism globally for the Cisco MDS switch. CHAP authentication is required for all iSCSI sessions. The validation is done using RADIUS or local authentication.

To configure the authentication policy for iSCSI sessions to a particular interface, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface GigabitEthernet 2/1.100 switch(config-if)#	Selects the Gigabit Ethernet interface.
Step 3	switch(config-if)# iscsi authentication none	Specifies that no authentication is required for iSCSI sessions to the selected interface.

The IPS module verifies the iSCSI host authentication using the local password database, TACACS+, or RADIUS (see the “[Configuring CLI User Profiles](#)” section on page 16-22). If local authentication is used, the **username iscsi-user password iscsi** command assigns a password and a user name for a new user. If the user name does not exist it will be created.

**Note**

The **iscsi** keyword is mandatory to identify iSCSI users.

To configure iSCSI users for local authentication, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# username iscsiuser password ffsffsfssfs345353554535 iscsi	Configures a user name (iscsiuser) and password (ffsffsfssfs345353554535) in the local database for iSCSI login authentication.
Note	The iscsi keyword is required at the end to identify the user.	

Advanced iSCSI Configuration

Advanced configuration options are available for iSCSI interfaces on a per-IPS port basis. These configurations are similar to the advanced FCIP configurations and are already explained in that section (see the [Advanced FCIP Profile Configuration](#), page 22-22).

To access these commands from the iSCSI interface, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# interface iscsi 4/1 switch(config-if)#	Selects the iSCSI interface on the switch that identifies all the initiators.
Step 3	switch(config-if)# tcp ? keepalive-timeout max-bandwidth-kbps max-bandwidth-mbps max-retransmissions min-retransmit-time pmtu-enable qos sack-enable send-buffer-size	Provides the TCP options available on a per-IPS port basis for iSCSI interfaces.

Cisco MDS switches support the following advanced features for iSCSI interfaces:

- iSCSI listener port—Configure the TCP port number for the iSCSI interface which listens for new TCP connections. The default port number is 3260. Following that, the iSCSI port only accepts TCP connections on the newly configured port (see the “[Configuring TCP Listener Ports](#)” section on page 22-22)
- TCP tuning parameters (see the “[Configuring TCP Parameters](#)” section on page 22-23). The following TCP parameters can be configured.
 - [Minimum Retransmit Timeout](#), page 22-24
 - [Keepalive Timeout](#), page 22-24
 - [Maximum Retransmissions](#), page 22-24
 - [Path MTU](#), page 22-25
 - [SACK](#), page 22-25—SACK is enabled by default for iSCSI TCP configurations.
 - [Window Management](#), page 22-25
 - [Buffer Size](#), page 22-26
 - [Monitoring Window Congestion](#), page 22-27
 - QoS—QoS configurations differ for iSCSI and FCIP interfaces. To set the QoS values, follow these steps:

	Command	Purpose
Step 1	<code>switch(config-if)# tcp qos 3</code>	Configures the control TCP connection. The DSCP value of 3 is applied to all IP packets in this iSCSI interface. The valid range for the iSCSI DSCP value is from 0 to 63.
Step 2	<code>switch(config-profile)# no tcp qos 5</code>	Reverts the switch to its factory default (marks all packets with DSCP value 0).

- Identification of dynamic iSCSI initiator—iSCSI initiators are identified based on their IQN name or their IP address. In the absence of any configuration for the initiator (WWN or VSAN membership), the identifier key is the default connection. By default, the key is the IQN name but can be changed to IP address by toggling this mode.
- Proxy or transparent Initiator—For each iSCSI initiator with iSCSI target sessions, the switch creates a virtual FC initiator with a distinct pair of WWNs per VSAN. Targets that have access control per LUN, the WWN pair of each FC initiator must be configured in the target. The proxy initiator mode can be enabled to facilitate this configuration, in which case, all iSCSI initiators which connect to this iSCSI interface inherit the same WWN pair and create only one virtual FC initiator in each VSAN.

iSCSI Forwarding Mode

The iSCSI gateway on the IPS module has two modes of forwarding operation

- The **pass-thru** mode (default): In this mode, the IPS port converts an iSCSI PDU into an FCP frame or vice versa and then forwards it one frame or PDU at a time. The absence of buffering PDUs or frames keeps the operation latency low. To operate in this mode, the IPS port has to negotiate with its peers a suitable maximum size of the data payload in each frame/PDU. This is done during iSCSI login and FC PLOGI and the value is restricted by the TCP connection's

Maximum Segment Size (MSS) and the maximum Fibre Channel data payload size specified by the FC target. This usually results in a smaller maximum payload size than most hosts expect, thus comes the second mode of forwarding.

- The **store-and-forward** mode: This mode allows the iSCSI client to send and receive an iSCSI data payload at the size it desires. This sometimes results in better performance for the client. The IPS port stores each TCP segment it receives until one full iSCSI PDU is received before converting and forwarding it as Fibre Channel frames to the FC target. In the opposite direction, the IPS port assembles all FC data frames of an exchange to build one iSCSI data-in PDU before forwarding it to the iSCSI client. The limitation on this mode is iSCSI CRC data digest cannot be used.

Displaying iSCSI Information

This section includes the following topics:

- [Displaying iSCSI Interfaces, page 22-64](#)
- [Displaying Global iSCSI Information, page 22-67](#)
- [Displaying iSCSI Sessions, page 22-67](#)
- [Displaying iSCSI Initiators, page 22-69](#)
- [Displaying iSCSI Virtual Targets, page 22-73](#)
- [Displaying IPS Statistics, page 22-73](#)
- [Displaying iSCSI User Information, page 22-75](#)

Displaying iSCSI Interfaces

Use the **show iscsi interface** command to view the summary, counter, description, and status of the iSCSI interface. Use the output to verify the administrative mode, the interface status, TCP parameters currently used, and brief statistics. See [Example 22-21](#).

Example 22-21 Displays the iSCSI Interface Information

```
switch# show interface iscsi 2/1
iscsi2/1 is up -----> Interface is up
  Hardware is GigabitEthernet
  Port WWN is 20:41:00:05:30:00:50:de
  Admin port mode is iSCSI
  Port mode is iSCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 7, Number of TCP connection: 7
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is disabled
    Keepalive-timeout is 1 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 8
    Sack is disabled
    Minimum available bandwidth is 0 kbps
    Estimated round trip time is 0 usec
  5 minutes input rate 265184 bits/sec, 33148 bytes/sec, 690 frames/sec
  5 minutes output rate 375002168 bits/sec, 46875271 bytes/sec, 33833 frames/sec
  iSCSI statistics
```

```

6202235 packets input, 299732864 bytes
  Command 6189718 pdus, Data-out 1937 pdus, 1983488 bytes, 0 fragments
146738794 packets output, 196613551108 bytes
  Response 6184282 pdus (with sense 4), R2T 547 pdus
  Data-in 140543388 pdus, 189570075420 bytes

```

The **show iscsi stats** command can be used to view brief or detailed iSCSI statistics per iSCSI interface. See Examples 22-22 and 22-23.

Example 22-22 Displays iSCSI Statistics for the iSCSI Interface

```

switch# show iscsi stats iscsi 4/1
iscsi4/1
  5 minutes input rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  iSCSI statistics
    1196 packets input, 173680 bytes
      Command 483 pdus, Data-out 104 pdus, 106496 bytes, 0 fragments
    output 1802 packets, 647152 bytes
      Response 483 pdus (with sense 0), R2T 25 pdus
      Data-in 685 pdus, 554696 bytes

```

Example 22-23 Displays Detailed iSCSI Statistics for the iSCSI Interface

```

switch# show iscsi stats iscsi 4/1 detail
iscsi4/1
  5 minutes input rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  iSCSI statistics
    1196 packets input, 173680 bytes
      Command 483 pdus, Data-out 104 pdus, 106496 bytes, 0 fragments
    output 1802 packets, 647152 bytes
      Response 483 pdus (with sense 0), R2T 25 pdus
      Data-in 685 pdus, 554696 bytes
  iSCSI Forward:
    Command: 483 PDUs (Rcvd: 483)
    Data-Out (Write): 104 PDUs (Rcvd 104), 0 fragments, 106496 bytes
  FCP Forward:
    Xfer_rdy: 25 (Rcvd: 25)
    Data-In: 685 (Rcvd: 719), 554696 bytes
    Response: 483 (Rcvd: 534), with sense 0
    TMF Resp: 0

  iSCSI Stats:
    Login: attempt: 25, succeed: 25, fail: 0, authen fail: 0
    Rcvd: NOP-Out: 556, Sent: NOP-In: 556
      NOP-In: 0, Sent: NOP-Out: 0
      TMF-REQ: 0, Sent: TMF-RESP: 0
      Text-REQ: 6, Sent: Text-RESP: 6
      SNACK: 0
      Unrecognized Opcode: 0, Bad header digest: 0
      Command in window but not next: 0, exceed wait queue limit: 0
      Received PDU in wrong phase: 0
      SCSI Busy responses: 0
  FCP Stats:
    Total: Sent: 726
      Received: 1366 (Error: 0, Unknown: 0)
    Sent: PLOGI: 17, Rcvd: PLOGI_ACC: 17, PLOGI_RJT: 0
      PRLI: 17, Rcvd: PRLI_ACC: 17, PRLI_RJT: 0, Error resp: 0
      LOGO: 12, Rcvd: LOGO_ACC: 0, LOGO_RJT: 0
      PRLO: 12, Rcvd: PRLO_ACC: 0, PRLO_RJT: 0

```

```

ABTS: 0, Rcvd: ABTS_ACC: 0
TMF REQ: 0
Self orig command: 51, Rcvd: data: 34, resp: 51
Rcvd: PLOGI: 20, Sent: PLOGI_ACC: 5, PLOGI_RJT: 15
LOGO: 5, Sent: LOGO_ACC: 5, LOGO_RJT: 0
PRLI: 5, Sent: PRLI_ACC: 5, PRLI_RJT: 0
PRLO: 0, Sent: PRLO_ACC: 0, PRLO_RJT: 0
ABTS: 0

iSCSI Drop:
Command: Target down 0, Task in progress 0, LUN map fail 0
      CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
      Persistent Resv 0, No task: 0
Data-Out: 0, Data CRC Error: 0
TMF-Req: 0, No task: 0
FCP Drop:
Xfer_rdy: 0, Data-In: 0, Response: 0

Buffer Stats:
Buffer less than header size: 0, Partial: 53, Split: 79
Pullup give new buf: 0, Out of contiguous buf: 0, Unaligned m_data: 0

```

Displaying Proxy Initiator Information

If the proxy initiator feature is enabled in the iSCSI interface, use the **show interface iscsi** command to display configured proxy initiator information:

Example 22-24 Displays Proxy Initiator Information for the iSCSI Interface with System-Assigned WWNs

```

switch# show interface iscsi 4/1
iscsi4/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:c1:00:05:30:00:a7:9e
  Admin port mode is iSCSI
  Port mode is iSCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0, Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is disabled
    QOS code point is 0
  Forwarding mode: pass-thru
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : enabled<-----Proxy initiator is enabled
    nWWN is 28:00:00:05:30:00:a7:a1 (system-assigned)<----System-assigned nWWN
    pWWN is 28:01:00:05:30:00:a7:a1 (system-assigned)<---- System-assigned pWWN
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  iSCSI statistics
    Input 7 packets, 2912 bytes
      Command 0 pdus, Data-out 0 pdus, 0 bytes
    Output 7 packets, 336 bytes
      Response 0 pdus (with sense 0), R2T 0 pdus
      Data-in 0 pdus, 0 bytes

```

Example 22-25 Displays Proxy Initiator Information for the iSCSI Interface with User-Assigned WWNs

```

switch# show interface iscsi 4/2
iscsi4/2 is up
  Hardware is GigabitEthernet
  Port WWN is 20:c1:00:05:30:00:a7:9e
  Admin port mode is ISCSI
  Port mode is ISCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0, Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is disabled
    QOS code point is 0
  Forwarding mode: pass-thru
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : enabled
    nWWN is 11:11:11:11:11:11:11:11 (manually-configured)<---user-assigned nWWN
    pWWN is 22:22:22:22:22:22:22:22 (manually-configured)<---user-assigned pWWN
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  iSCSI statistics
    Input 7 packets, 2912 bytes
      Command 0 pdus, Data-out 0 pdus, 0 bytes
    Output 7 packets, 336 bytes
      Response 0 pdus (with sense 0), R2T 0 pdus
      Data-in 0 pdus, 0 bytes

```

Displaying Global iSCSI Information

Use the **show iscsi global** command to view the overall configuration and the iSCSI status. See [Example 22-26](#)

Example 22-26 Displays the Current Global iSCSI Configuration and State.

```

switch# show iscsi global
iSCSI Global information
  Authentication: NONE
  Import FC Target: Enabled
  Number of target nodes: 5
  Number of portals: 8
  Number of sessions: 6
  Failed session: 0, Last failed initiator name:

```

Displaying iSCSI Sessions

Use the **show iscsi session** command to view details about the current iSCSI sessions in the switch. Without parameters, this command displays all sessions. The output can be filtered by specificity an initiator, a target, or both.

[Example 22-27](#) displays one iSCSI initiator configured based on the IQN name (iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k) and another based on it's IP address (10.10.100.199).

Example 22-27 Displays Brief Information of All iSCSI Sessions

```

switch# show iscsi session
Initiator ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  Session #1
    Discovery session, ISID 00023d000043, Status active

  Session #2
    Target VT1
    VSAN 1, ISID 00023d000046, Status active, no reservation

  Session #3
    Target VT2
    VSAN 1, ISID 00023d000048, Status active, no reservation

Initiator 10.10.100.199
  Initiator name ign.1987-05.com.cisco:01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT2
    VSAN 1, ISID 246700000000, Status active, no reservation

  Session #2
    Target VT1
    VSAN 1, ISID 246b00000000, Status active, no reservation

  Session #3
    Target ign.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
    VSAN 1, ISID 246e00000000, Status active, no reservation

```

[Example 22-28](#) and [Example 22-29](#) display the iSCSI initiator configured based on its IP address (10.10.100.199).

Example 22-28 Displays Brief Information About the Specified iSCSI Session

```

switch# show iscsi session initiator 10.10.100.199 target VT1
Initiator 10.10.100.199
  Initiator name ign.1987-05.com.cisco:01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT1
    VSAN 1, ISID 246b00000000, Status active, no reservation

```

Example 22-29 Displays Detailed Information About the Specified iSCSI Session

```

switch# show iscsi session initiator 10.10.100.199 target VT1 detail
Initiator 10.10.100.199 (oasis-qa)
  Initiator name ign.1987-05.com.cisco:01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1 (index 3)
    Target VT1
    VSAN 1, ISID 246b00000000, TSIH 384, Status active, no reservation
    Type Normal, ExpCmdSN 39, MaxCmdSN 54, Barrier 0
    MaxBurstSize 0, MaxConn 0, DataPDUIInOrder No
    DataSeqInOrder No, InitialR2T Yes, ImmediateData No
    Registered LUN 0, Mapped LUN 0
    Stats:
      PDU: Command: 38, Response: 38
      Bytes: TX: 8712, RX: 0
    Number of connection: 1
    Connection #1
      Local IP address: 10.10.100.200, Peer IP address: 10.10.100.199
      CID 0, State: LOGGED_IN

```



```

StatsN 62, ExpStatsN 0
MaxRecvDSLength 1024, our_MaxRecvDSLength 1392
CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
Version Min: 2, Max: 2
FC target: Up, Reorder PDU: No, Marker send: No (int 0)
Received MaxRecvDSLen key: No

```

Displaying iSCSI Initiators

Use the **show iscsi initiator** command to display information about all initiators connected to a iSCSI interface in the switch. The information can be filtered to display only the desired iSCSI initiator by specifying the initiator name. Detailed output of the iscsi initiator can be obtained by specifying the **detail** option. The **iscsi-session** (and optionally **detail**) parameter displays only iSCSI session information. The **fcp-session** (and optionally **detail**) parameter displays only FCP session information. The output includes static and dynamic initiators. See Examples 22-30 and 22-31.

Example 22-30 Displays Information About Connected iSCSI Initiators

```

switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco.02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1
  Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag: 0x180
    VSAN ID 1, FCID 0x6c0202
    VSAN ID 2, FCID 0x6e0000
    VSAN ID 10, FCID 0x790000

iSCSI Node name is 10.10.100.199
  iSCSI Initiator name: iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  iSCSI alias name: oasis-ga
  Node WWN is 22:03:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 5
  Number of Virtual n_ports: 1
  Virtual Port WWN is 22:00:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag: 0x180
    VSAN ID 5, FCID 0x640000
    VSAN ID 1, FCID 0x6c0203

```

Example 22-31 Display Detailed Information About the iSCSI Initiator

```

switch# show iscsi initiator iqn.1987-05.com.cisco.02.3021b0f2fda0.avanti12-w2k detail
iSCSI Node name is iqn.1987-05.com.cisco.02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1

  Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag is 0x180
    VSAN ID 1, FCID 0x6c0202
    1 FC sessions, 1 iSCSI sessions
    iSCSI session details          <-----iSCSI session details
      Target: VT1
      Statistics:

```

```

PDU: Command: 0, Response: 0
Bytes: TX: 0, RX: 0
Number of connection: 1
TCP parameters
  Local 10.10.100.200:3260, Remote 10.10.100.116:4190
  Path MTU: 1500 bytes
  Retransmission timeout: 310 ms
  Round trip time: Smoothed 160 ms, Variance: 38
  Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
  Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
  Congestion window: Current: 1 KB

FCP Session details          <-----FCP session details
Target FCID: 0x6c01e8 (S_ID of this session: 0x6c0202)
pWWN: 21:00:00:20:37:62:c0:0c, nWWN: 20:00:00:20:37:62:c0:0c
Session state: CLEANUP
1 iSCSI sessions share this FC session
  Target: VT1
Negotiated parameters
  RcvDataFieldSize 1392 our_RcvDataFieldSize 1392
  MaxBurstSize 0, EMPD: FALSE
  Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
  PDU: Command: 0, Response: 0

```

Use the **show fcns database** (and optionally **detail**) to display the Fibre Channel name server entry for the Fibre Channel N port created for iSCSI initiators in the SAN. See Examples 22-32 to 22-35.

Example 22-32 Displays the Fibre Channel Name Server Database

```

switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x020101      N      22:04:00:05:30:00:35:e1 (Cisco)          scsi-fcp:init isc..w <--iSCSI
0x020102      N      22:02:00:05:30:00:35:e1 (Cisco)          scsi-fcp:init isc..w initiator
0x0205d4      NL     21:00:00:04:cf:da:fe:c6 (Seagate)        scsi-fcp:target
0x0205d5      NL     21:00:00:04:cf:e6:e4:4b (Seagate)        scsi-fcp:target
0x0205d6      NL     21:00:00:04:cf:e6:21:ac (Seagate)        scsi-fcp:target
0x0205d9      NL     21:00:00:04:cf:e6:19:9b (Seagate)        scsi-fcp:target
0x0205da      NL     21:00:00:04:cf:e6:19:62 (Seagate)        scsi-fcp:target
0x0205dc      NL     21:00:00:04:cf:e6:e9:82 (Seagate)        scsi-fcp:target
0x0205e0      NL     21:00:00:04:cf:e6:21:06 (Seagate)        scsi-fcp:target
0x0205e1      NL     21:00:00:04:cf:e6:e0:eb (Seagate)        scsi-fcp:target

Total number of entries = 10

VSAN 2:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0xef0001      N      22:02:00:05:30:00:35:e1 (Cisco)          scsi-fcp:init isc..w

Total number of entries = 1

VSAN 3:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0xed0001      N      22:02:00:05:30:00:35:e1 (Cisco)          scsi-fcp:init isc..w

Total number of entries = 1

```

Example 22-33 Displays the FCNS Database in Detail

```

switch# show fcns database detail
-----
VSAN:1      FCID:0x020101
-----
port-wwn (vendor)      :22:04:00:05:30:00:35:e1 (Cisco)
node-wwn                :22:03:00:05:30:00:35:e1
class                   :2,3
node-ip-addr            :10.2.2.12
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name      :
symbolic-node-name      :iqn.1991-05.com.microsoft:oasis2-dell
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :22:01:00:05:30:00:35:de
hard-addr               :0x000000
-----
VSAN:1      FCID:0x020102
-----
port-wwn (vendor)      :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn                :22:01:00:05:30:00:35:e1
class                   :2,3
node-ip-addr            :10.2.2.11
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name      :
symbolic-node-name      :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :22:01:00:05:30:00:35:de
hard-addr               :0x000000
-----
...
Total number of entries = 10
=====
-----
VSAN:2      FCID:0xef0001
-----
port-wwn (vendor)      :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn                :22:01:00:05:30:00:35:e1
class                   :2,3
node-ip-addr            :10.2.2.11
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name      :
symbolic-node-name      :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :22:01:00:05:30:00:35:de
hard-addr               :0x000000
-----
Total number of entries = 1
-----
...

```

Example 22-34 Displays Detailed Information for a Fibre Channel N Port Created for An iSCSI Initiator Identified by it's IQN Name

```
switch# show fcns database fcid 0x6c0203 detail vsan 1
-----
VSAN:1      FCID:0x6c0203
-----
port-wwn (vendor)      :22:00:00:05:30:00:10:e1 (Cisco)
node-wwn               :22:03:00:05:30:00:10:e1
class                 :2,3
node-ip-addr           :10.10.100.199
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :10.10.100.199
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :20:c1:00:05:30:00:10:de
hard-addr              :0x000000

Total number of entries = 1
```

Example 22-35 Displays Detailed Information for a Fibre Channel N Port created for An iSCSI Initiator Identified by it's IP Address

```
switch# show fcns database fcid 0x6c0203 detail vsan 1
-----
VSAN:1      FCID:0x6c0203
-----
port-wwn (vendor)      :22:00:00:05:30:00:10:e1 (Cisco)
node-wwn               :22:03:00:05:30:00:10:e1
class                 :2,3
node-ip-addr           :10.10.100.199
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :10.10.100.199<----ID assigned by IP address
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :20:c1:00:05:30:00:10:de
hard-addr              :0x000000

Total number of entries = 1
```

Use the **show iscsi initiator configured** to display information about all the configured iSCSI initiators. Specifying the name shows information about the desired initiator. See [Example 22-36](#).

Example 22-36 Display Information About Configured Initiators

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Member of vsans: 1, 2, 10
  Node WWN is 22:01:00:05:30:00:10:e1
  No. of PWWN: 5
    Port WWN is 22:04:00:05:30:00:10:e1
    Port WWN is 22:05:00:05:30:00:10:e1
    Port WWN is 22:06:00:05:30:00:10:e1
    Port WWN is 22:07:00:05:30:00:10:e1
    Port WWN is 22:08:00:05:30:00:10:e1

iSCSI Node name is 10.10.100.199
  Member of vsans: 1, 5
```

```

Node WWN is 22:03:00:05:30:00:10:e1
No. of PWWN: 4
  Port WWN is 22:00:00:05:30:00:10:e1
  Port WWN is 22:09:00:05:30:00:10:e1
  Port WWN is 22:0a:00:05:30:00:10:e1
  Port WWN is 22:0b:00:05:30:00:10:e1

```

Displaying iSCSI Virtual Targets

Use the **show iscsi virtual-target** to display information about the FC targets exported as iSCSI virtual targets to the iSCSI initiators. The output includes static as well as dynamic targets. See [Example 22-37](#).

Example 22-37 Displays Exported Targets

```

switch# show iscsi virtual-target
target: VT1
  * Port WWN 21:00:00:20:37:62:c0:0c
    Configured node
    all initiator permit is enabled

target: VT2
  Port WWN 21:00:00:04:cf:4c:52:c1
  Configured node
  all initiator permit is disabled

target: ign.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
  Port WWN 21:00:00:20:37:a6:be:32 , VSAN 1
  Auto-created node

```

Displaying IPS Statistics

The **show ips stats tcp interface** command displays information about the underlying transport for iSCSI. See Examples [22-38](#) and [22-39](#).

Example 22-38 Displays iSCSI Stats (brief)

```

switch# show ips stats tcp interface gigabitethernet 2/1
TCP Statistics for port GigabitEthernet2/1
  Connection Stats
    0 active openings, 6 accepts
    0 failed attempts, 0 reset received, 6 established
  Segment stats
    640780835 received, 150953931 sent, 12 retransmitted
    0 bad segments received, 0 reset sent
  TCP Active Connections

```

Local	Address	Remote Address	State	Send-Q	Recv-Q
10.48.69.250:3260		10.48.69.226:1026	ESTABLISH	0	0
10.48.69.250:3260		10.48.69.231:1026	ESTABLISH	0	0
10.48.69.250:3260		10.48.69.231:1033	ESTABLISH	0	0
10.48.69.250:3260		10.48.69.226:1038	ESTABLISH	0	0
0.0.0.0:3260		0.0.0.0:0	LISTEN	0	0

Example 22-39 Displays SCSI Stats (detail)

```

switch# show ips stats tcp interface gigabitethernet 2/1 detail
TCP Statistics for port GigabitEthernet2/1
  TCP send stats
    150953931 segments, 2755572300 bytes
    53986369 data, 82341597 ack only packets

```

```

4 control (SYN/FIN/RST), 0 probes, 14625949 window updates
12 segments retransmitted, 576 bytes
12 retransmitted while on ethernet send queue, 0 packets split
118741734 delayed acks sent
TCP receive stats
640780835 segments, 640325552 data packets in sequence, 925034009772 bytes in
sequence
0 predicted ack, 615117910 predicted data
0 bad checksum, 0 multi/broadcast, 0 bad offset
0 no memory drops, 0 short segments
0 duplicate bytes, 0 duplicate packets
0 partial duplicate bytes, 0 partial duplicate packets
0 out-of-order bytes, 0 out-of-order packets
0 packet after window, 0 bytes after window
0 packets after close
25656078 acks, 2755572210 ack bytes, 0 ack toomuch, 5786 duplicate acks
0 ack packets left of snd_una, 0 non-4 byte aligned packets
12100 window updates, 0 window probe
29 pcb hash miss, 17 no port, 0 bad SYN, 0 paws drops
TCP Connection Stats
0 attempts, 6 accepts, 6 established
4 closed, 4 drops, 0 conn drops
0 drop in retransmit timeout, 4 drop in keepalive timeout
0 drop in persist drops, 0 connections drained
TCP Miscellaneous Stats
21635776 segments timed, 21642712 rtt updated
12 retransmit timeout, 0 persist timeout
8494 keepalive timeout, 8490 keepalive probes
TCP SACK Stats
0 recovery episodes, 0 data packets, 0 data bytes
0 data packets retransmitted, 0 data bytes retransmitted
0 connections closed, 0 retransmit timeouts
TCP SYN Cache Stats
6 entries, 6 connections completed, 0 entries timed out
0 dropped due to overflow, 0 dropped due to RST
0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
0 hash collisions, 0 retransmitted
TCP Active Connections

```

Local Address	Remote Address	State	Send-Q	Recv-Q
10.48.69.250:3260	10.48.69.226:1026	ESTABLISH	0	0
10.48.69.250:3260	10.48.69.231:1026	ESTABLISH	0	0
10.48.69.250:3260	10.48.69.231:1033	ESTABLISH	0	0
10.48.69.250:3260	10.48.69.226:1038	ESTABLISH	0	0
0.0.0.0:3260	0.0.0.0:0	LISTEN	0	0

The **show ips stats buffer** command displays information about the iSCSI buffers. See Example 22-40

Example 22-40 Displays iSCSI Buffers

```

switch# show ips stats buffer interface gigabitethernet 4/2
Mbuf Statistics for port GigabitEthernet4/2
Free Mbufs : 83221
Mbuf high watermark : 124830
Mbuf low watermark : 20805
Mbuf alloc failures : 0
Total clusters : 2304
Free Clusters : 80145
Clusters high watermark : 87381
Clusters low watermark : 79059
Clusters alloc failures : 0
Free shared mbufs : 0
Shared Mbuf alloc failures : 0

```

```
Free shared clusters          : 0
Shared clusters alloc failures: 0

Ether channel Statistics for port GigabitEthernet4/2
TCP segments sent            : 0
TCP segments received        : 0
Xmit packets sent            : 0
Xmit packets received        : 0
Config packets sent          : 0
Config packets received      : 0
MPQ packets send errors      : 0
```

Displaying iSCSI User Information

The **show user-account iscsi** command displays all configured iSCSI user names. See [Example 22-41](#).

Example 22-41 Displays iSCSI User Names

```
switch# show user-account iscsi
username:iscsiuser
secret: dsfffsffsffasffsdffg

username:user2
secret:cshadhdhsadadjajdjas
```

iSCSI High Availability

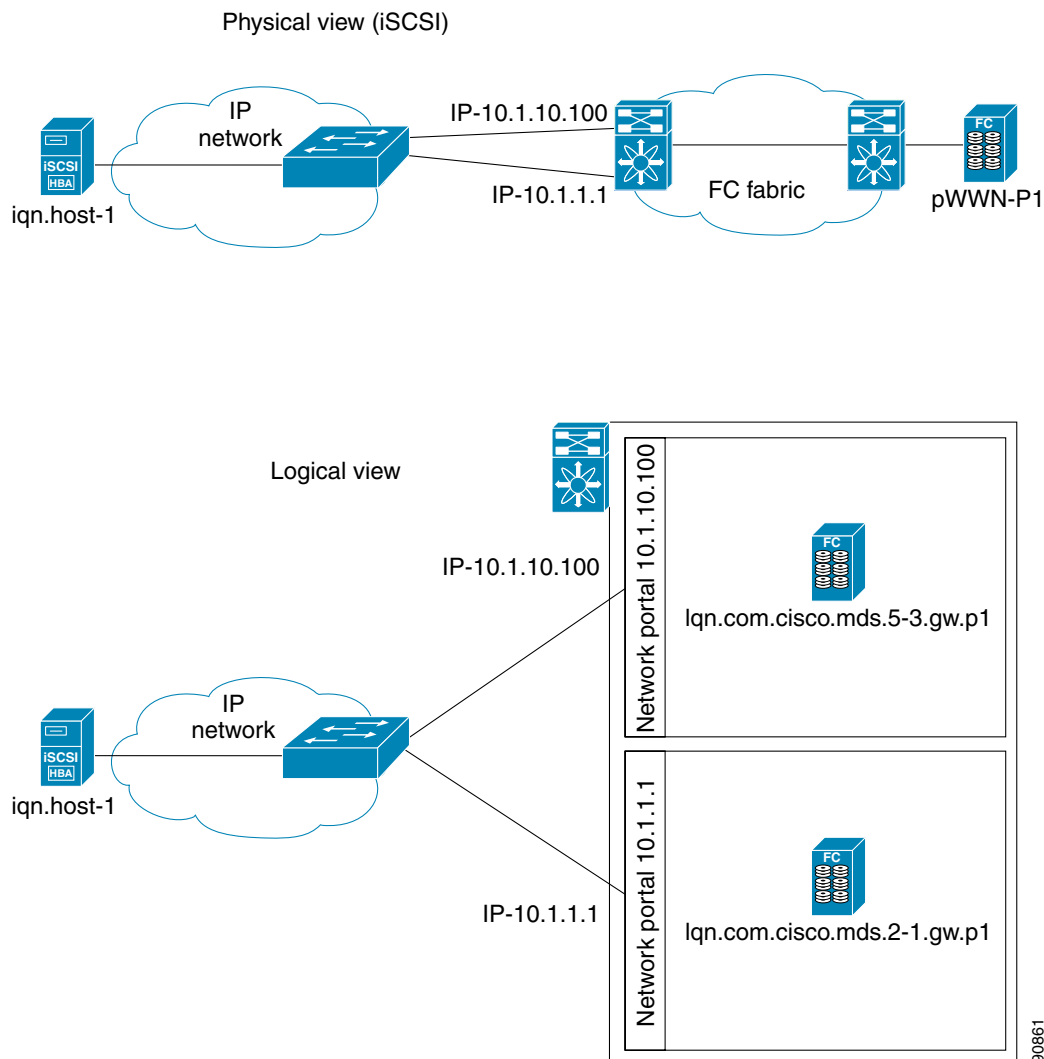
The following high availability features are available for iSCSI configurations:

- [Multiple IPS Ports Connected to the Same IP Network, page 22-76](#)
- [VRRP-Based High Availability, page 22-77](#)
- [Ethernet PortChannel-Based High Availability, page 22-78](#)

Multiple IPS Ports Connected to the Same IP Network

Figure 22-31 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Figure 22-31 Multiple Gigabit Ethernet Interfaces in the Same IP Network

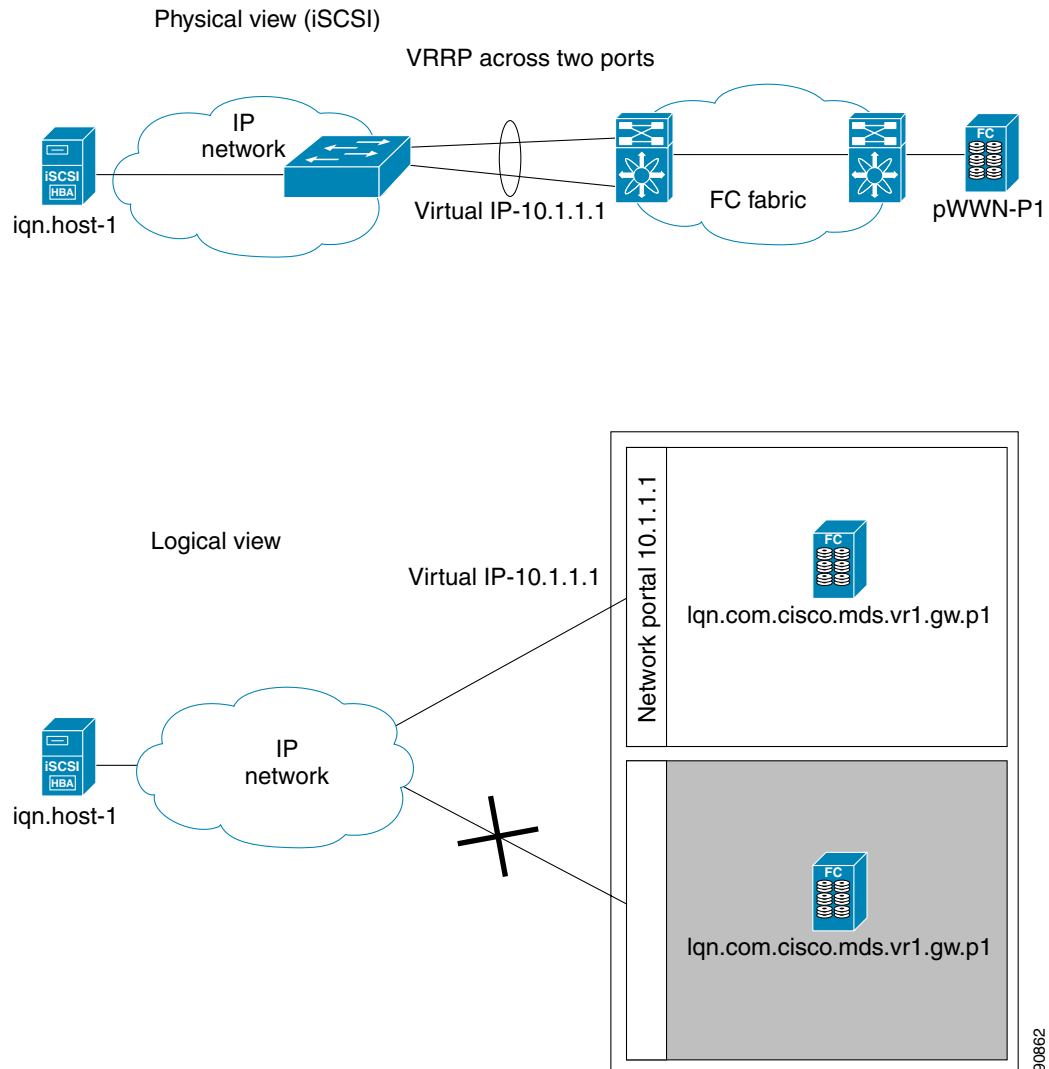


In Figure 22-31, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

VRRP-Based High Availability

Figure 22-32 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 22-32 VRRP-Based iSCSI High Availability



In Figure 22-32, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.



Tip

Ports that act as VRRP master and backup can be on different switches. If you have a static WWN configuration for iSCSI initiators (see the [“Presenting iSCSI Hosts as Virtual Fibre Channel Hosts”](#) section on page 22-52), configure a different WWN for the iSCSI initiator for each switch. If you use a proxy-initiator, be sure to configure a different pWWN on each iSCSI interface for each VRRP port used.

Ethernet PortChannel-Based High Availability

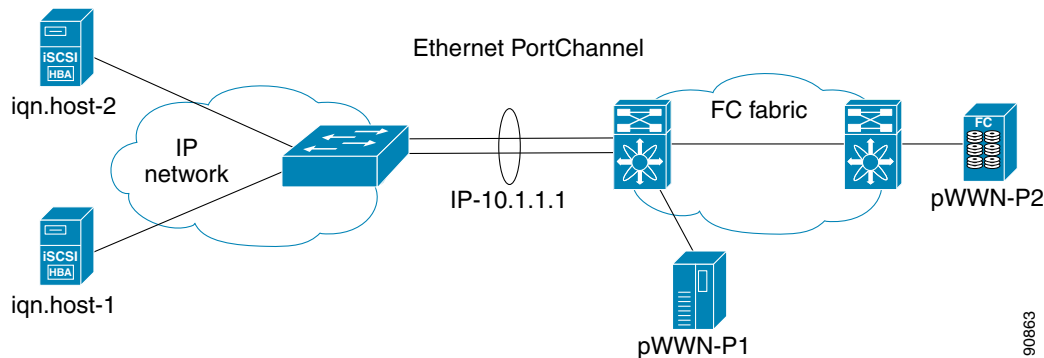


Note

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth will be one Gbps for that iSCSI link.

Figure 22-33 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 22-33 Ethernet PortChannel-Based iSCSI High Availability



In Figure 22-33, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the virtual iSCSI target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.

iSCSI Authentication Setup Guidelines

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios. It includes the following authentication setup guidelines:

- [No Authentication, page 22-78](#)
- [CHAP with Local Password Database, page 22-79](#)
- [CHAP with External RADIUS Server, page 22-79](#)



Note

This section does not specify the steps to enter or exit EXEC mode, configuration mode, or any submode. Be sure to verify the prompt before issuing any command.

No Authentication

To configure a network with no authentication set the iSCSI authentication method to none.

```
switch(config)# iscsi authentication none
```

CHAP with Local Password Database

To configure authentication using the CHAP option with the local password database, follow these steps:

- Step 1** Set the AAA authentication to use the local password database for iSCSI protocol.

```
switch(config)# aaa authentication iscsi default local
```

- Step 2** Set the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 3** Configure the user names and passwords for iSCSI users.

```
switch(config)# username iscsi-user password abcd iscsi
```



Note If you do not specify the **iscsi** option, the user name is assumed to be a MDS switch login user instead of an iSCSI user.

- Step 4** Verify the global iSCSI authentication setup.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----verify
  Import FC Target: Disabled
  ...
```

CHAP with External RADIUS Server

To configure authentication using the CHAP option with an external RADIUS server, follow these steps:

- Step 1** Configure the password for the MDS as RADIUS client to the RADIUS server.

```
switch(config)# radius-server key mds-1
```

- Step 2** Configure the RADIUS server IP address.

```
switch(config)# radius-server host 10.1.1.10
```

- Step 3** Configure a server group

```
switch(config)# aaa group server radius iscsi-radius-group
switch(config-radius)# server 10.1.1.1
```

- Step 4** Setup the authentication verification for iSCSI protocol to go to RADIUS server.

```
switch(config)# aaa authentication iscsi default group iscsi-radius-group
```

- Step 5** Setup the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 6** Verify that the global iSCSI authentication setup is CHAP.

```
switch# show iscsi global
iSCSI Global information
  Authentication: CHAP          <----- Verify CHAP
  ....
```

Step 7 Verify that the aaa authentication information for iSCSI

```

switch# show aaa authentication
      default: local
      console: local
      iscsi: group iscsi-radius-group    <----- Group name
      dhchap: local

switch# show radius-server groups
total number of groups:2

following RADIUS server groups are configured:
  group radius:
    server: all configured radius servers
  group iscsi-radius-group:
    server: 10.1.1.1 on auth-port 1812, acct-port 1813

switch# show radius-server
Global RADIUS shared secret:mds-1    <----- Verify secret
....

following RADIUS servers are configured:
  10.1.1.1:    <----- Verify the server IP address
    available for authentication on port:1812
    available for accounting on port:1813

```

To configure an iSCSI RADIUS server, follow these steps:

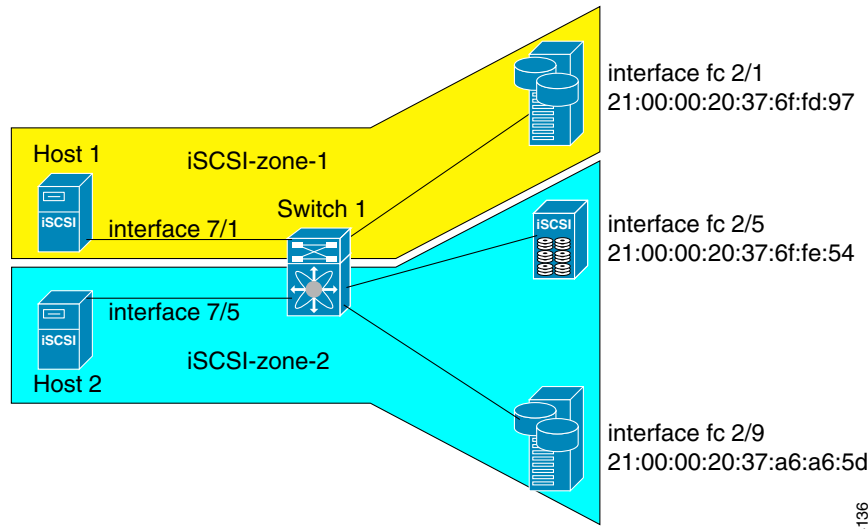
-
- Step 1** Configure the RADIUS server to allow access from the MDS switch's management Ethernet IP address.
 - Step 2** Configure the shared secret for the RADIUS server to authenticate the MDS switch.
 - Step 3** Configure the iSCSI users and passwords on the RADIUS server.
-

Scenario 1

Sample scenario 1 assumes the following configuration (see [Figure 22-34](#)):


- Access control using Fibre Channel zoning.
- No target-based LUN mapping or LUN masking.
- No iSCSI authentication (none).
- iSCSI initiator identified using IP address (Host 1 = 10.11.1.10).
- iSCSI initiator identified using node name (Host 2 = iqn.1987-05.com.cisco:01.25589167f74c).

Figure 22-34 iSCSI Scenario 1



94136

To configure scenario 1 (see [Figure 22-34](#)), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts.
- ```
switch(config)# iscsi authentication none
```
- Step 2** Configure iSCSI to dynamically import all FC targets into the iSCSI SAN using auto-generated iSCSI target names.
- ```
switch(config)# iscsi import target fc
```
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```
-  **Note** Host 1 is connected to this port.
- 
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by the IP address, and enable the interface.
- ```
switch(config)# int iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```
- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name, and enable the interface.
- ```
switch(config)# int iscsi 7/5
switch(config-if)# switchport initiator id name
```

```
switch(config-if)# no shut
```



Note Host 1 is connected to this port.

Step 7 Verify the available Fibre Channel targets (see [Figure 22-34](#)).

```
switch# show fcns database
VSAN 1:
```

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x6d0001	NL	21:00:00:20:37:6f:fd:97	(Seagate)	scsi-fcp:target
0x6d0101	NL	21:00:00:20:37:6f:fe:54	(Seagate)	scsi-fcp:target
0x6d0201	NL	21:00:00:20:37:a6:a6:5d	(Seagate)	scsi-fcp:target

Total number of entries = 3

Step 8 Create a zone named *iscsi-zone-1* with host 1 and one Fibre Channel target in it.



Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

```
switch(config)# zone name iscsi-zone-1 vsan 1
switch(config-zone)# member pwn 21:00:00:20:37:6f:fd:97
switch(config-zone)# member symbolic-nodename 10.11.1.10
```

Step 9 Create a zone named *iscsi-zone-2* with host 2 and two FC targets in it.



Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

```
switch(config)# zone name iscsi-zone-2 vsan 1
switch(config-zone)# member pwn 21:00:00:20:37:6f:fe:54
switch(config-zone)# member pwn 21:00:00:20:37:a6:a6:5d
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c
```

Step 10 Create a zoneset and add the two zones as members.

```
switch(config)# zoneset name zoneset-iscsi vsan 1
switch(config-zoneset)# member iscsi-zone-1
switch(config-zoneset)# member iscsi-zone-2
```

Step 11 Activate the zoneset.

```
switch(config)# zoneset activate name zoneset-iscsi vsan 1
```

Step 12 Display the active zoneset.



Note The iSCSI hosts has not connected so they do not have a FCID yet.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwn 21:00:00:20:37:6f:fd:97] <-----Target
      symbolic-nodename 10.11.1.10 <-----iSCSI host (host 1, not online)

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwn 21:00:00:20:37:6f:fe:54] <-----Target
```

```
* fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d] <-----Target
symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c <-iSCSI host (host 2, not online)
```

Step 13 Bring up the iSCSI hosts (Host 1 and Host 2).

Step 14 Show all the iSCSI sessions (use the **detail** option for detailed information).

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.25589167f74c <-----Host 2
Initiator ip addr (s): 10.15.1.11
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54
```



Note The last part of the auto-created target name is the FC target's pWWN.

```
VSAN 1, ISID 00023d000001, Status active, no reservation

Session #2
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d
VSAN 1, ISID 00023d000001, Status active, no reservation

Initiator 10.11.1.10 <-----Host 1
Initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97
VSAN 1, ISID 00023d000001, Status active, no reservation
```

Step 15 Verify the details of the two iSCSI initiators.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.25589167f74c <-----
Initiator ip addr (s): 10.15.1.11
iSCSI alias name: oasis11.cisco.com
Node WWN is 20:02:00:0b:fd:44:68:c2 (dynamic)
Member of vsans: 1
Number of Virtual n_ports: 1
Virtual Port WWN is 20:03:00:0b:fd:44:68:c2 (dynamic)
Interface iSCSI 7/5, Portal group tag: 0x304
VSAN ID 1, FCID 0x6d0300

iSCSI Node name is 10.11.1.10 <-----
iSCSI Initiator name: iqn.1987 - 05.com.cisco:01.e41695d16b1a
iSCSI alias name: oasis10.cisco.com
Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic)
Member of vsans: 1
Number of Virtual n_ports: 1
Virtual Port WWN is 20:05:00:0b:fd:44:68:c2 (dynamic)
Interface iSCSI 7/1, Portal group tag: 0x300
VSAN ID 1, FCID 0x6d0301
```

Host 2: Initiator ID based on node name because the initiator is entering iSCSI interface 7/5

Host 1: Initiator ID based on IP address because the initiator is entering iSCSI interface 7/1

Step 16 View the active zone set. The iSCSI initiators' FCIDs are resolved.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97]
    * fcid 0x6d0301 [symbolic-nodename 10.11.1.10] <-----

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d]
    * fcid 0x6d0300 [symbolic-nodename
iqn.1987-05.com.cisco:01.25589167f74c] <-----
```

**FCID resolved for
Host 1**

FCID for Host 2

Step 17 The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                      (VENDOR)      FC4-TYPE:FEATURE
-----
0x6d0001      NL    21:00:00:20:37:6f:fd:97  (Seagate)     scsi-fcp:target
0x6d0101      NL    21:00:00:20:37:6f:fe:54  (Seagate)     scsi-fcp:target
0x6d0201      NL    21:00:00:20:37:a6:a6:5d  (Seagate)     scsi-fcp:target
0x6d0300      N     20:03:00:0b:fd:44:68:c2  (Cisco)       scsi-fcp:init isc..w
0x6d0301      N     20:05:00:0b:fd:44:68:c2  (Cisco)       scsi-fcp:init isc..w
```


Step 18 Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

```
switch# show fcns database fcid 0x6d0300 detail vsan 1
-----
VSAN:1      FCID:0x6d0300
-----
port-wwn (vendor)      :20:03:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:02:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr            :10.15.1.11  <-----
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name      :

symbolic-node-name
:iqn.1987-05.com.cisco:01.25589167f74c<-----
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :21:91:00:0b:fd:44:68:c0
hard-addr               :0x000000
Total number of entries = 1

switch# show fcns database fcid 0x6d0301 detail vsan 1
-----
VSAN:1      FCID:0x6d0301
-----
port-wwn (vendor)      :20:05:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:04:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr            :10.11.1.10
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name      :

symbolic-node-name      :10.11.1.10  <-----
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :21:81:00:0b:fd:44:68:c0
hard-addr               :0x000000
```

IP address of the iSCSI
host

iSCSI gateway node

iSCSI initiator ID is
based on the registered
node name

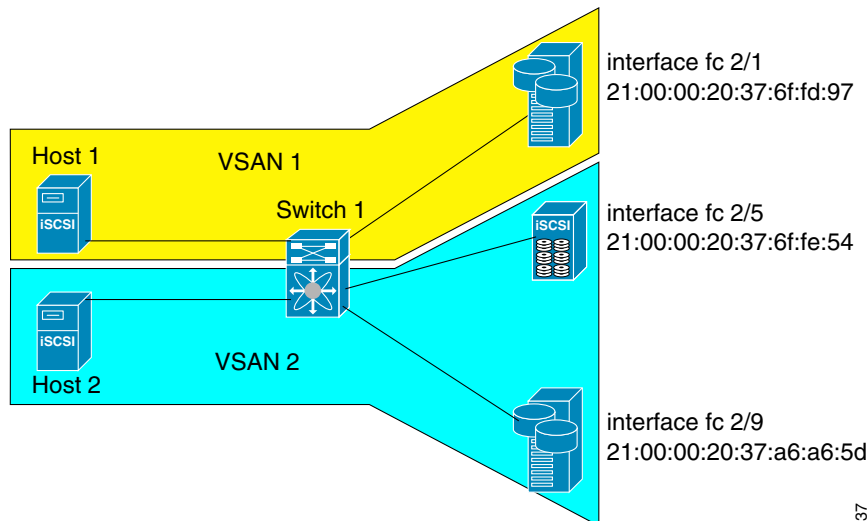
iSCSI gateway node

iSCSI initiator ID is
based on the IP address
registered in
symbolic-node-name
field

Scenario 2

Sample scenario 2 (see [Figure 22-35](#)) assumes the following configuration:

- Access control based on Fibre Channel zoning.
- Target-based LUN mapping or LUN masking.
- No iSCSI authentication (none).
- iSCSI initiator assigned to different VSANs.

Figure 22-35 iSCSI Scenario 2

94137

To configure scenario 1 (see [Figure 22-35](#)), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts.
- ```
switch(config)# iscsi authentication none
```
- Step 2** Configure iSCSI to dynamically import all FC targets into the iSCSI SAN using auto-generated iSCSI target names.
- ```
switch(config)# iscsi import target fc
```
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iscsi initiators by the IP address, and enable the interface.
- ```
switch(config)# int iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```
- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```
- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iscsi initiators by IP address, and enable the interface.
- ```
switch(config)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```
- Step 7** Add static configuration for each iSCSI initiator.
- ```
switch(config)# iscsi initiator name iqn.1987-05.com.cisco:01.e41695d16b1a <-----Host 2
```

```
switch(config-(iscsi-init))# static pWWN system-assign 1
switch(config-(iscsi-init))# static nWWN system-assign

switch(config)# iscsi initiator ip address 10.15.1.11
switch(config-(iscsi-init))# static pwwn system-assigned 1
switch(config-(iscsi-init))# vsan 2
```

**Step 8** View the configured initiators.



**Note** The WWNs are assigned by the system. The initiators are members of different VSANs.

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
 Member of vsans: 1
 Node WWN is 20:03:00:0b:fd:44:68:c2
 No. of PWWN: 1
 Port WWN is 20:02:00:0b:fd:44:68:c2

iSCSI Node name is 10.15.1.11
 Member of vsans: 2
 No. of PWWN: 1
 Port WWN is 20:06:00:0b:fd:44:68:c2
```

**Step 9** Create a zone with Host 1.

```
switch(config)# zone name iscsi-zone-1 vsan 1
```

**Step 10** Add three members to the zone named *iscsi-zone-1*.



**Note** Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN can be used. In this case, the pWWN is persistent.

- Based on the symbolic node name.

```
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.e41695d16b1a
```

- Based on the persistent pWWN assigned to the initiator. You can obtain the pWWN from the **show iscsi initiator** output.

```
switch(config-zone)# member pwwn 20:02:00:0b:fd:44:68:c2
```

**Step 11** Create zone with Host 2 and two Fibre Channel targets.



**Note** If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

```
switch(config)# zone name iscsi-zone-2 vsan 2
```

**Step 12** Activate the zoneset in VSAN 2

```
switch(config)# zoneset activate name iscsi-zoneset-v2 vsan 2
Zoneset activation initiated. check zone status
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
 zone name iscsi-zone-2 vsan 2
 * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
 * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]
 pwwn 20:06:00:0b:fd:44:68:c2 <-----Host is not online
```

**Step 13** Start the iSCSI clients on both hosts and verify that sessions come up.

**Step 14** Display the iSCSI sessions to verify the Fibre Channel target and the configured WWNs.

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.e41695d16b1a
 Initiator ip addr (s): 10.11.1.10
 Session #1
 Discovery session, ISID 00023d000001, Status active

 Session #2
 Target
iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97<---- To FC target
 VSAN 1, ISID 00023d000001, Status active, no reservation
```

**Step 15** Display the iSCSI initiator to verify the configured nWWN and pWWN.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
 Initiator ip addr (s): 10.11.1.10
 iSCSI alias name: oasis10.cisco.com

 Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<----- The configured nWWN
 Member of vsans: 1
 Number of Virtual n_ports: 1

 Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<---- The configured pWWN
 Interface iSCSI 7/1, Portal group tag: 0x300
 VSAN ID 1, FCID 0x680102
```

**Step 16** Check the Fibre Channel name server

```
switch# show fcns database vsan 1
VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x680001 NL 21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target

0x680102 N 20:02:00:0b:fd:44:68:c2 (Cisco) scsi-fcp:init isc..w<--- iSCSI initiator in
 name server
```

**Step 17** Verify the details of the iSCSI initiator's FCID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1

VSAN:1 FCID:0x680102

port-wwn (vendor) :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn :20:03:00:0b:fd:44:68:c2
class :2,3
node-ip-addr :10.11.1.10
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :21:81:00:0b:fd:44:68:c0
 iSCSI alias name: oasis10.cisco.com

Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<-----
Member of vsans: 1
Number of Virtual n_ports: 1
```

**The configured nWWN**

```
Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<---
Interface iSCSI 7/1, Portal group tag: 0x300
VSAN ID 1, FCID 0x680102
```

**The configured pWWN****Step 18** Check the Fibre Channel name server

```
switch# show fcns database vsan 1
VSAN 1:

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x680001 NL 21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target
0x680102 N 20:02:00:0b:fd:44:68:c2 (Cisco) scsi-fcp:init isc..w <-----
```

**iSCSI  
initiator in  
name server****Step 19** Verify the details of the iSCSI initiator's FCID in the name server

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1

VSAN:1 FCID:0x680102

port-wwn (vendor) :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn :20:03:00:0b:fd:44:68:c2
class :2,3
node-ip-addr :10.11.1.10
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :21:81:00:0b:fd:44:68:c0
hard-addr :0x000000
```

**Step 20** Verify that zoning has resolved the FCID for the iSCSI client.

```
switch# show zoneset active vsan 1
zoneset name iscsi-zoneset-v1 vsan 1
zone name iscsi-zone-1 vsan 1
* fcid 0x680001 [pwwn 21:00:00:20:37:6f:fd:97]
* fcid 0x680102 [pwwn 20:02:00:0b:fd:44:68:c2]
```

**Step 21** Do the same to verify that the second initiator is connected to the two Fibre Channel targets in VSAN 2.

```
switch# show iscsi session initiator 10.15.1.11
Initiator 10.15.1.11
Initiator name ign.1987-05.com.cisco:01.25589167f74c
Session #1
Target ign.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54 <--
VSAN 2, ISID 00023d000001, Status active, no reservation

Session #2
Target ign.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d <--
VSAN 2, ISID 00023d000001, Status active, no reservation
```

**Session to  
first target**

**Session to  
second  
target**

```
switch# show iscsi initiator
iSCSI Node name is 10.15.1.11 <--- Initiator ID is the IP address
iSCSI Initiator name: ign.1987-05.com.cisco:01.25589167f74c
iSCSI alias name: oasis11.cisco.com
```

```
Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic) <-----
Member of vsans: 2 <--- vsan membership
Number of Virtual n_ports: 1
```

**Dynamic  
WWN as  
static WWN  
not  
assigned**

```
Virtual Port WWN is 20:06:00:0b:fd:44:68:c2 (configured) <-----
Interface iSCSI 7/5, Portal group tag: 0x304
VSAN ID 2, FCID 0x750200
```

**Static  
pWWN for  
the initiator**

```
switch# show fcns database vsan 2
VSAN 2:
```

```

FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE

0x750001 NL 21:00:00:20:37:6f:fe:54 (Seagate) scsi-fcp:target
0x750101 NL 21:00:00:20:37:a6:a6:5d (Seagate) scsi-fcp:target

0x750200 N 20:06:00:0b:fd:44:68:c2 (Cisco) scsi-fcp:init isc..w <--
Total number of entries = 3
```

**iSCSI  
initiator  
entry in  
name server**

```

switch# show fcns database fcid 0x750200 detail vsan 2

VSAN:2 FCID:0x750200

port-wwn (vendor) :20:06:00:0b:fd:44:68:c2 (Cisco)
node-wwn :20:04:00:0b:fd:44:68:c2
class :2,3
node-ip-addr :10.15.1.11
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name :
symbolic-node-name :10.15.1.11
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :21:91:00:0b:fd:44:68:c0
hard-addr :0x000000
Total number of entries = 1

switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
 zone name iscsi-zone-2 vsan 2
 * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
 * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]

 * fcid 0x750200 [pwwn 20:06:00:0b:fd:44:68:c2] <-----

```

**FCID  
resolved for  
iSCSI  
initiator**

## Configuring Storage Name Services

Effective Release 1.3(1), the Internet Storage Name Service (iSNS) client feature is available in all switches in the Cisco MDS 9000 Family with IPS modules installed.

iSNS services allow your existing TCP/IP networks to function more effectively as storage area networks by automating the discover and management of iSCSI devices. To facilitate these functions, the iSNS client functionality registers iSCSI portals and all targets accessible through a particular interface with an external iSNS server.

## Creating iSNS Profiles and Tagging Profiles

The iSNS client functionality on each interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with its configured iSNS server using an iSNS profile. This process is referred to as tagging an iSNS profile to an interface. Each iSNS profile keeps information about an iSNS server IP address. One profile can be tagged to one or more interfaces.

Once a profile is tagged to an interface, the MDS switch opens a TCP connection to the iSNS server IP address (using a well-known iSNS port number 3205) in the profile and registers network entity and portal objects. It goes through the FC name server database and configuration to find storage nodes to register with the server.

Statically-mapped virtual target (see the “[Presenting Fibre Channel Targets as iSCSI Targets](#)” section on [page 22-46](#)) is registered if the associated Fibre channel pWWN is present in the FC name server database and no access control configuration prevents it (using the **advertise interface** or the **initiator** options in the **iscsi virtual-target name** command). A dynamically-mapped target is registered if the dynamic target importing is enabled using the **iscsi import target fc** command.

A storage node is deregistered from the iSNS server when it becomes unavailable either because of configuration change (such as access control change or dynamic import disabling) or when the Fibre Channel storage port goes off-line. It will be registered again when the node is online.

When the iSNS client is unable to register/deregister objects with the iSNS server (e.g. as in unable to make a tcp connection to the iSNS server) it retries every minute to re-registers all iSNS objects for the affected interface(s) with the iSNS server.

Untagging a profile causes the network entity and portal to deregister from that interface.

To create an iSNS profile, follow these steps:

|        | Command                                                                           | Purpose                                               |
|--------|-----------------------------------------------------------------------------------|-------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                        | Enters configuration mode.                            |
| Step 2 | switch(config)# <b>isns profile name MyIsns</b><br>switch(config-(isns-profile))# | Creates a profile called MyIsns.                      |
|        | switch(config)# <b>no isns profile name OldIsns</b>                               | Removes a configured iSNS profile called OldIsns.     |
| Step 3 | switch(config-(isns-profile))# <b>server 10.10.100.211</b>                        | Specifies an iSNS server IP address for this profile. |
|        | switch(config-(isns-profile))# <b>no server 10.20.100.211</b>                     | Removes a configured iSNS server from this profile.   |

To modify the iSNS profile:

- for an interface, untag the interface from currently tagged profile and then tag to a new profile
- for a profile, remove the existing server and then add the new server.

To tag an interface to a profile, follow these steps:

|        | Command                                                                    | Purpose                                              |
|--------|----------------------------------------------------------------------------|------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                 | Enters configuration mode.                           |
| Step 2 | switch(config)# <b>interface gigabitethernet 4/1</b><br>switch(config-if)# | Configures the specified Gigabit Ethernet interface. |
| Step 3 | switch(config-if)# <b>isns MyIsns</b>                                      | Tags this interface to the profile.                  |

To untag an interface from a profile, follow these steps:

|        | Command                                                                    | Purpose                                              |
|--------|----------------------------------------------------------------------------|------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                 | Enters configuration mode.                           |
| Step 2 | switch(config)# <b>interface gigabitethernet 5/1</b><br>switch(config-if)# | Configures the specified Gigabit Ethernet interface. |
| Step 3 | switch(config-if)# <b>no isns OldIsns</b>                                  | Untags this interface from the profile.              |



All associated iSNS objects for an interface tagged to an iSNS profile can be re-registered with the iSNS server using the **isns reregister** command in EXEC mode.

```
switch# isns reregister gigabitethernet 1/4
switch# isns reregister port-channel 1
switch#
```

## Verifying iSNS Configurations

Use the **show isns profile** command to view configured iSNS profiles. Profile ABC has two portals registered with the iSNS server. Each portal corresponds to a particular interface. Profile XYZ has a specified iSNS server, but does not have any tagged interfaces configured (see [Example 22-42](#)).

### Example 22-42 Displays all Configured iSNS Profiles

```
switch# show isns profile
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204

iSNS profile name XYZ
iSNS Server 10.10.100.211
```

### Example 22-43 Displays a Specified iSNS Profile

```
switch# show isns profile ABC

iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204
```

Use the **show isns profile counters** command to view all configured profiles with the iSNS PDU statistics for each tagged interface.

### Example 22-44 Displays Configured Profiles with iSNS Statistics

```
switch# show isns profile counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
 Input 54 pdus (registration/deregistration pdus only)
 Reg pdus 37, Dereg pdus 17
 Output 54 pdus (registration/deregistration pdus only)
 Reg pdus 37, Dereg pdus 17
iSNS Server 10.10.100.204

iSNS profile name XYZ
tagged interface port-channel 2
iSNS statistics
 Input 30 pdus (registration/deregistration pdus only)
 Reg pdus 29, Dereg pdus 1
 Output 30 pdus (registration/deregistration pdus only)
 Reg pdus 29, Dereg pdus 1
iSNS Server 10.1.4.218
```

**Example 22-45 Displays a Specified Profile's iSNS Statistics**

```

switch# show isns profile ABC counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
 Input 54 pdus (registration/deregistration pdus only)
 Reg pdus 37, Dereg pdus 17
 Output 54 pdus (registration/deregistration pdus only)
 Reg pdus 37, Dereg pdus 17
iSNS Server 10.10.100.204

```

Use the show isns query to view all objects registered on the iSNS server and specified in the given profile.

**Example 22-46 Displays iSNS Queries**

```

switch# show isns query ABC gigabitethernet 2/3
iSNS server: 10.10.100.204
Init: iqn.1991-05.com.microsoft:ibmw2k
 Alias: <MS SW iSCSI Initiator>
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03.210000203762fa34
 nWWN: 200000203762fa34

```

Use the **show interface** command to view the iSNS profile to which a interface is tagged (see [Example 22-47](#)).

**Example 22-47 Displays Tagged iSNS Interfaces**

```

switch# show int gigabitethernet 2/3
GigabitEthernet2/3 is up
Hardware is GigabitEthernet, address is 0005.3000.ae94
Internet address is 10.10.100.201/24
MTU 1500 bytes
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
Auto-Negotiation is turned on
iSNS profile ABC

5 minutes input rate 112 bits/sec, 14 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1935 packets input, 132567 bytes
 4 multicast frames, 0 compressed
 0 input errors, 0 frame, 0 overrun 0 fifo
1 packets output, 42 bytes, 0 underruns
 0 output errors, 0 collisions, 0 fifo
 0 carrier errors

```

# Default IP Storage Settings

Table 22-2 lists the default settings for Gigabit Ethernet parameters.

**Table 22-2 Default Gigabit Ethernet Parameters**

| Parameters        | Default                           |
|-------------------|-----------------------------------|
| IP MTU frame size | 1500 bytes for all Ethernet ports |

Table 22-3 lists the default settings for FCIP parameters.

**Table 22-3 Default FCIP Parameters**

| Parameters                              | Default                     |
|-----------------------------------------|-----------------------------|
| TCP default port for FCIP               | 3225                        |
| minimum-retransmit-time                 | 200 milliseconds.           |
| keepalive-timeout                       | 60 seconds.                 |
| max-retransmissions                     | 4 retransmissions.          |
| PMTU discovery                          | Enabled.                    |
| pmtu-enable reset-timeout               | 3600 seconds.               |
| SACK                                    | Enabled.                    |
| max-bandwidth                           | 1G.                         |
| min-available-bandwidth                 | 15 Mbps.                    |
| round-trip-time                         | 1 ms.                       |
| buffer size                             | 0 KB.                       |
| Control TCP and data connection         | No packets are transmitted. |
| TCP congestion window monitoring        | Enabled                     |
| Burst size                              | 10KB.                       |
| TCP connection mode                     | active mode is enabled.     |
| special-frame                           | Disabled.                   |
| FCIP timestamp                          | Disabled.                   |
| acceptable-diff range to accept packets | + or - 1000 milliseconds.   |
| B port keepalive responses              | Disabled                    |

Table 22-4 lists the default settings for iSCSI parameters.

**Table 22-4 Default iSCSI Parameters**

| Parameters                     | Default                |
|--------------------------------|------------------------|
| Number of TCP connections      | One per iSCSI session. |
| Fibre Channel targets to iSCSI | Not imported.          |

**Table 22-4 Default iSCSI Parameters (continued)**

| Parameters                                         | Default                                                                                                             |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Advertising iSCSI target                           | Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces |
| iSCSI hosts mapping to virtual Fibre Channel hosts | Dynamic mapping.                                                                                                    |
| Dynamic iSCSI initiators                           | Members of the VSAN 1.                                                                                              |
| Identifying initiators                             | iSCSI node names.                                                                                                   |
| Advertising static virtual targets                 | No initiators allowed to access a virtual target (unless explicitly configured).                                    |
| iSCSI login authentication                         | CHAP or none authentication mechanism.                                                                              |
| Ethernet PortChannel IP address usage              | Source and destination IP addresses.                                                                                |



## Configuring Call Home

---

Call Home provides e-mail-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications. Common uses of this feature may include direct paging of a network support engineer, e-mail notification to a Network Operations Center, and utilization of Cisco AutoNotify services for direct case generation with the Technical Assistance Center.

This chapter provides configuration and messaging details on the Call Home feature. It includes the following sections:

- [Call Home Features, page 23-2](#)
- [Call Home Configuration Process, page 23-2](#)
- [Cisco AutoNotify, page 23-3](#)
- [Configuring the Call Home Function, page 23-3](#)
- [Assigning Contact Information, page 23-4](#)
- [Configuring Destination Profiles, page 23-5](#)
- [Configuring Alert Groups, page 23-7](#)
- [Configuring Message Levels, page 23-8](#)
- [Configuring E-Mail Options, page 23-9](#)
- [Enabling or Disabling Call Home, page 23-9](#)
- [Testing Call Home Communication, page 23-10](#)
- [Displaying Call Home Information, page 23-10](#)
- [Default Settings, page 23-12](#)
- [Event Triggers, page 23-13](#)
- [Call Home Message Severity Levels, page 23-15](#)
- [Message Contents, page 23-16](#)

# Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Family. It provides multiple Call Home profiles (also referred to as Call Home destination profiles), each with separate potential destinations. Each profile may be predefined or user-defined.

The Call Home function can even leverage support from Cisco Systems or another support partner. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts and trigger events on the switch.
- Automatic execution and attachment of relevant command output.
- Multiple message format options:
  - Short Text—Suitable for pagers or printed reports.
  - Plain Text—Full formatted message information suitable for human reading.
  - XML—Matching readable format using Extensible Markup Language (XML) and Document Type Definitions (DTDs) named Messaging Markup Language (MML). The MML DTD is published on the Cisco Connection Online (CCO) website at <http://www.cisco.com/>. The XML format enables communication with the Cisco Systems TAC group.
- Multiple concurrent message destinations. Up to 50 E-mail destination addresses are allowed for each format type.
- Message categories include system, environment, switching module hardware, supervisor module, hardware, inventory, and test.

## Call Home Configuration Process

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

- E-mail server and at least one destination profile (predefined or user-defined) must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, email, or automated service such as Cisco AutoNotify.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This is required to determine the origin of messages received.
- The Cisco MDS 9000 switch must have IP connectivity to an E-mail server for the feature to operate.
- If Cisco AutoNotify is used, an active service contract must cover the device being configured.

To configure Call Home, follow these steps:

- 
- |               |                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Configure the Call Home function (see the “ <a href="#">Configuring the Call Home Function</a> ” section on page 23-3). |
| <b>Step 2</b> | Assign contact information (see the “ <a href="#">Assigning Contact Information</a> ” section on page 23-4).            |
| <b>Step 3</b> | Configure destination profiles (see the “ <a href="#">Configuring Destination Profiles</a> ” section on page 23-5).     |
| <b>Step 4</b> | Enable or disable Call Home (see the “ <a href="#">Enabling or Disabling Call Home</a> ” section on page 23-9).         |
| <b>Step 5</b> | Test Call Home messages (see the “ <a href="#">Testing Call Home Communication</a> ” section on page 23-10).            |
-

# Cisco AutoNotify

For those who have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible through registration with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support. To register, the following items are required:

- The SMARTnet contract number covering your MDS 9000 family switch.
- Your name, company address, your email address, and your CCO ID.
- The serial number of your Cisco MDS 9000 Family switch. This can be obtained by looking at the serial number label on the back of the switch (next to the power supply), or by executing the operating system **show sprom backplane 1** command.
- The exact product number of your Cisco MDS 9000 Family switch. This can be obtained by executing the same operating system command as above. For example, some valid product numbers include: DS-C6509 and DS-C9216-K9

To configure a Cisco MDS 9000 Family switch to use AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and email address information is found on the Cisco.com web site at:

[http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtton/anoti\\_ds.htm](http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtton/anoti_ds.htm)



## Note

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, email server, and an XML destination profile as specified in the Service Activation document ([http://www.cisco.com/univercd/cc/td/doc/product/voice/c\\_callmg/3\\_3/service/serv332/ccm/srvs/sssrvtact.htm](http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_3/service/serv332/ccm/srvs/sssrvtact.htm)). The **contract-id**, **customer-id**, **site-id**, and **switch-priority** parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

## Configuring the Call Home Function

To enter the Call Home submode, follow these steps:

|        | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Purpose                                        |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Enters configuration mode.                     |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Enters Call Home submode.                      |
| Step 3 | switch(config-callhome)# <b>?</b><br>contract-id      Service contract id of the customer<br>customer-id      Customer id<br>destination-profile    Configure destination profiles<br>disable            Disable callhome<br>email-contact        Email address of the contact person<br>enable              Enable callhome<br>exit                Exit from this submode<br>no                  Negate a command or set its defaults<br>phone-contact        Contact person's phone number<br>site-id              Site id of the network where switch is deployed<br>streetaddress        Configure replacement part shipping address.<br>switch-priority      Priority of the switch(0-highest 7-lowest)<br>transport            Configure transport related configuration | Displays the options available at this prompt. |

# Assigning Contact Information

It is mandatory for each switch to include e-mail, phone, and street address information. It's optional to include the contract ID, customer ID, site ID, and switch priority information.

To assign the contact information, follow these steps:

|         | Command                                                                                                                                                              | Purpose                                                                                                                                                                                             |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1  | switch# <b>config t</b>                                                                                                                                              | Enters configuration mode.                                                                                                                                                                          |
| Step 2  | switch# <b>snmp-server contact</b><br>personname@companyname.com                                                                                                     | Configures the SNMP contact e-mail address to receive a test message reply from Cisco.                                                                                                              |
| Step 3  | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                                          | Enters the Call Home submode.                                                                                                                                                                       |
| Step 4  | switch(config-callhome)# <b>email-contact</b><br>username@company.com<br>successfully updated the information<br>switch(config-callhome)#                            | Assigns the customer's e-mail address. Up to 128 alphanumeric characters are accepted in e-mail address format.<br><br><b>Note</b> You can use any valid e-mail address. You cannot use spaces.     |
| Step 5  | switch(config-callhome)# <b>phone-contact</b><br>+1-800-123-4567<br>successfully updated the information<br>switch(config-callhome)#                                 | Assigns the customer's phone number. Up to 20 alphanumeric characters are accepted in international format.<br><br><b>Note</b> You cannot use spaces. Be sure to use the + prefix before the number |
| Step 6  | switch(config-callhome)# <b>streetaddress</b> 1234<br>Picaboo Street, Any city, Any state, 12345<br>successfully updated the information<br>switch(config-callhome)# | Assigns the customer's street address where the equipment is located. Up to 256 alphanumeric characters are accepted in free format.                                                                |
| Step 7  | switch(config-callhome)# <b>switch-priority</b> 0<br>successfully updated the information<br>switch(config-callhome)#                                                | Assigns the switch priority, with 0 being the highest priority and 7 the lowest.<br><br><b>Tip</b> Use this field to create a hierarchical management structure.                                    |
| Step 8  | switch(config-callhome)# <b>customer-id</b><br>Customer1234<br>successfully updated the information<br>switch(config-callhome)#                                      | Optional. Identifies the customer ID. Up to 256 alphanumeric characters are accepted in free format.                                                                                                |
| Step 9  | switch(config-callhome)# <b>site-id</b><br>Site1ManhattanNY<br>successfully updated the information<br>switch(config-callhome)#                                      | Optional. Identifies the customer site ID. Up to 256 alphanumeric characters are accepted in free format.                                                                                           |
| Step 10 | switch(config-callhome)# <b>contract-id</b><br>Company1234<br>successfully updated the information<br>switch(config-callhome)#                                       | Assigns the customer ID for the switch. Up to 64 alphanumeric characters are accepted in free format.                                                                                               |



# Configuring Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator. At least one destination profile is required. You can configure multiple destination profiles of one or more types.

You can use one of the predefined destination profiles or define a desired profile. If you define a new profile, you must assign a profile name.



## Note

If you use the Cisco AutoNotify service, the XML destination profile is required (see [http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtton/anoti\\_ds.htm](http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtton/anoti_ds.htm)).

- **Profile Name**—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are **full-txt**, **short-txt**, or **XML** (default).
- **Destination address**—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- **Message formatting**—The message format used for sending the alert (full text, short text, or XML).

To configure predefined destination profile messaging options, follow these steps:

|        | Command                                                                                                                      | Purpose                                                                                                                                                                                                                                                         |
|--------|------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                      | Enters configuration mode.                                                                                                                                                                                                                                      |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                  | Enters the Call Home submode.                                                                                                                                                                                                                                   |
| Step 3 | switch(config-callhome)#<br><b>destination-profile</b><br><b>full-txt-destination email-addr</b><br><b>person@place.com</b>  | Configures a predefined destination e-mail address for a message sent in full text format. This text provides the complete, detailed explanation of the failure.<br><br><b>Tip</b> Use a standard e-mail address that does not have any text size restrictions. |
|        | switch(config-callhome)#<br><b>destination-profile</b><br><b>full-txt-destination message-size</b><br><b>1000000</b>         | Configures a predefined destination message size for a message sent in full text format. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.                                       |
|        | switch(config-callhome)#<br><b>destination-profile</b><br><b>short-txt-destination email-addr</b><br><b>person@place.com</b> | Configures a predefined destination e-mail address for a message sent in short text format. This text provides the basic explanation of the failure.<br><br><b>Tip</b> Use a pager-related e-mail address for this option.                                      |
|        | switch(config-callhome)#<br><b>destination-profile</b><br><b>short-txt-destination message-size</b><br><b>100000</b>         | Configures a predefined destination message size for a message sent in short text format. The valid range is 0 to 1,000,000 bytes and the default is 4000. A value of 0 implies that a message of any size can be sent.                                         |

| Command                                                                                                               | Purpose                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 5</b><br><pre>switch(config-callhome)# destination-profile XML-destination email-addr findout@cisco.com</pre> | Configures a predefined destination e-mail address for a message sent in XML format. This option provides the full information that is compatible with Cisco Systems TAC support.<br><br><b>Tip</b> Do not add a pager-related e-mail address to this destination profile because of the large message size. |
| <pre>switch(config-callhome)# destination-profile XML-destination message-size 100000</pre>                           | Configures a predefined destination message size for a message sent in XML format. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.                                                                                          |

**Note**

Steps 3, 4, and 5 in this procedure can be skipped or configured in any order.

To configure new destination profile messaging options, follow these steps:

| Command                                                                                                   | Purpose                                                                                                                                                                                                                                |
|-----------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b><br><pre>switch# config t</pre>                                                              | Enters configuration mode.                                                                                                                                                                                                             |
| <b>Step 2</b><br><pre>switch(config)# callhome switch(config-callhome)#</pre>                             | Enters the Call Home submenu.                                                                                                                                                                                                          |
| <b>Step 3</b><br><pre>switch(config-callhome)# destination-profile test</pre>                             | Configures a new destination profile called test.                                                                                                                                                                                      |
| <b>Step 4</b><br><pre>switch(config-callhome)# destination-profile test email-addr person@place.com</pre> | Configures the e-mail address for the user-defined destination message (test) sent in default XML format.                                                                                                                              |
| <pre>switch(config-callhome)# destination-profile test message-size 1000000</pre>                         | Configures a message size for the user-defined destination message (test) sent in default XML format. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent. |
| <b>Step 5</b><br><pre>switch(config-callhome)# destination-profile test format full-txt</pre>             | Configures a user-defined destination message (test) sent in full text format.                                                                                                                                                         |
| <pre>switch(config-callhome)# destination-profile test format short -txt</pre>                            | Configures a user-defined destination message (test) sent in short text format.                                                                                                                                                        |

**Note**

Steps 4 and 5 in this procedure can be skipped or configured in any order.

# Configuring Alert Groups

The **alert-group** option allows you to select predefined types of Call Home alert notifications for destination profiles (predefined and user-defined). Destination profiles can be associated with multiple alert groups.

To configure alert group options, follow these steps:

|        | Command                                                                                                   | Purpose                                                                                                                                                |
|--------|-----------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                                             | Enters configuration mode.                                                                                                                             |
| Step 2 | <code>switch(config)# callhome</code><br><code>switch(config-callhome)#</code>                            | Enters Call Home submode.                                                                                                                              |
| Step 3 | <code>switch(config-callhome)# destination-profile test1 alert-group test</code>                          | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for all user-generated test events.              |
|        | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group test</code>          | Optional. Configures predefined short-text destination message profile to send Call Home notifications for all user-generated test events.             |
| Step 4 | <code>switch(config-callhome)# destination-profile test1 alert-group all</code>                           | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for all events.                                  |
|        | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group all</code>           | Optional. Configures predefined short-text destination message profile to send Call Home notifications for all events                                  |
| Step 5 | <code>switch(config-callhome)# destination-profile test1 alert-group Cisco-TAC</code>                     | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for events which are meant only for Cisco TAC.   |
|        | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group Cisco-TAC</code>     | Optional. Configures predefined short-text destination message profile to send Call Home notifications for events which are meant only for Cisco TAC.  |
| Step 6 | <code>switch(config-callhome)# destination-profile test1 alert-group environmental</code>                 | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for power, fan, and temperature-related events.  |
|        | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group environmental</code> | Optional. Configures predefined short-text destination message profile to send Call Home notifications for power, fan, and temperature-related events. |
| Step 7 | <code>switch(config-callhome)# destination-profile test1 alert-group inventory</code>                     | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for inventory status events.                     |
|        | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group inventory</code>     | Optional. Configures predefined short-text destination message profile to send Call Home notifications for inventory status events.                    |

|         | Command                                                                                                         | Purpose                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Step 8  | <code>switch(config-callhome)# destination-profile test1 alert-group linecard-hardware</code>                   | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for module-related events.      |
|         | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group linecard-hardware</code>   | Optional. Configures predefined short-text destination message profile to send Call Home notifications for module-related events.     |
| Step 9  | <code>switch(config-callhome)# destination-profile test1 alert-group supervisor-hardware</code>                 | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for supervisor-related events.  |
|         | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group supervisor-hardware</code> | Optional. Configures predefined short-text destination message profile to send Call Home notifications for supervisor-related events. |
| Step 10 | <code>switch(config-callhome)# destination-profile test1 alert-group system</code>                              | Optional. Configures user-defined destination message profile (test1) to send Call Home notifications for software-related events.    |
|         | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group system</code>              | Optional. Configures predefined short-text destination message profile to send Call Home notifications for software-related events.   |

## Configuring Message Levels

The **message-level** option allows you to filter messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower than the urgency threshold will not be sent. The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages will be sent).

To configure alert group options, follow these steps:

|        | Command                                                                              | Purpose                                                                                                       |
|--------|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                        | Enters configuration mode.                                                                                    |
| Step 2 | <code>switch(config)# callhome</code><br><code>switch(config-callhome)#</code>       | Enters Call Home submode.                                                                                     |
| Step 3 | <code>switch(config-callhome)# destination-profile test message-level 5</code>       | Optional. Configures the message level urgency as 5 and above for the user-defined profile (test1).           |
|        | <code>switch(config-callhome)# no destination-profile oldtest message-level 7</code> | Removes a previously configured urgency level and reverts it to the default of 0 (all messages will be sent). |

## Configuring E-Mail Options

You can configure the from, reply-to, and return-receipt e-mail addresses. While most e-mail address configurations are optional, you must ensure to configure the SMTP server address and port number for the Call Home functionality to work.

### Configuring General E-Mail Options

To configure general e-mail options, follow these steps:

|        | Command                                                                   | Purpose                                                                                 |
|--------|---------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                   | Enters configuration mode.                                                              |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#               | Enters Call Home submode.                                                               |
| Step 3 | switch(config-callhome)# <b>transport email from user@company1.com</b>    | Optional. Configures the from e-mail address.                                           |
| Step 4 | switch(config-callhome)# <b>transport email reply-to person@place.com</b> | Optional. Configures the reply-to e-mail address to which all responses should be sent. |

### Configuring SMTP Server and Ports

To configure the SMTP server and port, follow these steps:

|        | Command                                                                                                                                             | Purpose                                                                                                                         |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>                                                                                                                             | Enters configuration mode.                                                                                                      |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)#                                                                                         | Enters Call Home submode.                                                                                                       |
| Step 3 | switch(config-callhome)# <b>transport email smtp-server 192.168.1.1</b><br>successfully updated the information<br>switch(config-callhome)#         | Configures the DNS or IP address of the SMTP server to reach the server. The port usage defaults to 25 if no port is specified. |
|        | switch(config-callhome)# <b>transport email smtp-server 192.168.1.1 port 30</b><br>successfully updated the information<br>switch(config-callhome)# | <b>Note</b> The port number is optional and, if required, may be changed depending on the server location.                      |

## Enabling or Disabling Call Home

Once you have configured the contact information, you must enable the Call Home function. The **enable** command is required for the Call Home function to start operating.

To enable the Call Home function, follow these steps:

|        | Command                                                     | Purpose                    |
|--------|-------------------------------------------------------------|----------------------------|
| Step 1 | switch# <b>config t</b>                                     | Enters configuration mode. |
| Step 2 | switch(config)# <b>callhome</b><br>switch(config-callhome)# | Enters Call Home submode.  |

|        | Command                                                                                             | Purpose                                                                                                                                                                                                                 |
|--------|-----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 3 | switch(config-callhome)# <b>enable</b><br>callhome enabled successfully<br>switch(config-callhome)# | Enables the Call Home function.                                                                                                                                                                                         |
|        | switch(config-callhome)# <b>disable</b><br>switch(config-callhome)#                                 | Disables the Call Home function. When you disable the Call Home function, all input events are ignored.<br><br><b>Note</b> Even if Call Home is disabled, basic information for each Call Home event is sent to syslog. |

## Testing Call Home Communication

You can simulate a message generation by issuing a **test** command.

To test the Call Home function, follow these steps:

|        | Command                                                                                                                              | Purpose                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| Step 1 | switch# <b>callhome test</b><br>trying to send test callhome message<br>successfully sent test callhome message<br>switch#           | Sends a test message to the configured destination(s).           |
| Step 2 | switch# <b>callhome test inventory</b><br>trying to send test callhome message<br>successfully sent test callhome message<br>switch# | Sends a test inventory message to the configured destination(s). |

## Displaying Call Home Information

Use the **show callhome** command to display the configured Call Home information (see Examples 23-1 to 23-7).

### Example 23-1 Displays Configured Call Home Information

```
switch# show callhome
callhome enabled
Callhome Information:
contact person name:who@where
contact person's email:person@place.com
contact person's phone number:310-408-4000
street addr:1234 Picaboo Street, Any city, Any state, 12345
site id:Site1ManhattanNewYork
customer id:Customer1234
contract id:Andiamo1234
switch priority:0
```

### Example 23-2 Displays Information for All Destination Profiles (Predefined and User-defined)

```
switch# show callhome destination-profile
XML destination profile information
maximum message size:500000
message format:XML
message-level:0
email addresses configured:
alert groups configured:
```

```
cisco_tac

test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
cchetty@isco.com

alert groups configured:
test

full-txt destination profile information
maximum message size:500000
message format:full-txt
message-level:0
email addresses configured:

alert groups configured:
all

short-txt destination profile information
maximum message size:4000
message format:short-txt
message-level:0
email addresses configured:

alert groups configured:
all
```

### **Example 23-3** *Displays Information for a User-defined Destination Profile*

```
switch# show callhome destination-profile test
test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
user@company.com

alert groups configured:
test
```

### **Example 23-4** *Displays the Full-Text Profile*

```
switch# show callhome destination-profile profile full-txt-destination
full-txt destination profile information
maximum message size:250000
email addresses configured:
person2@company2.com
```

### **Example 23-5** *Displays the Short-Text Profile*

```
switch# show callhome destination-profile profile short-txt-destination
Short-txt destination profile information
maximum message size:4000
email addresses configured:
person2@company2.com
```

**Example 23-6 Displays the XML Destination Profile**

```
switch# show callhome destination-profile profile XML-destination
XML destination profile information
maximum message size:250000
email addresses configured:
findout@cisco.com
```

**Example 23-7 Displays E-mail and SMTP Information**

```
switch# show callhome transport-email
from email addr:user@company1.com
reply to email addr:pointer@company.com
return receipt email addr:user@company1.com
smtp server:server.company.com
smtp server port:25
```

## Default Settings

[Table 23-1](#) lists the default Call Home default settings.

**Table 23-1 Default Call Home Settings**

| Parameters                                                                        | Default |
|-----------------------------------------------------------------------------------|---------|
| Destination message size for a message sent in full text format.                  | 500,000 |
| Destination message size for a message sent in XML format.                        | 500,000 |
| Destination message size for a message sent in short text format.                 | 4,000   |
| DNS or IP address of the SMTP server to reach the server if no port is specified. | 25      |



# Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned commands to execute when the event occurs. The command output is included in the transmitted message. [Table 23-2](#) lists the trigger events. [Table 23-3](#) lists event categories and command outputs.

**Table 23-2 Event Triggers**

| Event     | Alert Group                                            | Event Name                   | Description                                                                                      | Severity Level |
|-----------|--------------------------------------------------------|------------------------------|--------------------------------------------------------------------------------------------------|----------------|
| Call Home | System and CISCO_TAC                                   | SW_CRASH                     | A software process has crashed with a stateless restart, indicating an interruption of a service | 5              |
|           | System and CISCO_TAC                                   | SW_SYSTEM_INCONSISTENT       | Inconsistency detected in software or file system                                                | 5              |
|           | Environmental and CISCO_TAC                            | TEMPERATURE_ALARM            | Thermal sensor indicates temperature reached operating threshold.                                | 6              |
|           |                                                        | POWER_SUPPLY_FAILURE         | Power supply failed.                                                                             | 6              |
|           |                                                        | FAN_FAILURE                  | Cooling fan has failed.                                                                          | 5              |
|           | Switching module and CISCO_TAC                         | LINECARD_FAILURE             | Switching module operation failed.                                                               | 7              |
|           |                                                        | POWER_UP_DIAGNOSTICS_FAILURE | Switching module failed power up diagnostics.                                                    | 7              |
|           | Line Card Hardware and CISCO_TAC                       | PORT_FAILURE                 | Hardware failure of interface port(s)                                                            | 6              |
|           | Line Card Hardware, Supervisor Hardware, and CISCO_TAC | BOOTFLASH_FAILURE            | Failure of boot compact flash card                                                               | 6              |
|           | Supervisor module and CISCO_TAC                        | SUP_FAILURE                  | Supervisor module operation failed.                                                              | 7              |
|           |                                                        | POWER_UP_DIAGNOSTICS_FAILURE | Supervisor module failed power up diagnostics.                                                   | 7              |

**Table 23-2 Event Triggers (continued)**

| Event     | Alert Group                       | Event Name         | Description                                                                                               | Severity Level |
|-----------|-----------------------------------|--------------------|-----------------------------------------------------------------------------------------------------------|----------------|
| Call Home | Supervisor Hardware and CISCO_TAC | INBAND_FAILURE     | Failure of inband communications path                                                                     | 7              |
|           | Supervisor Hardware and CISCO_TAC | EOBC_FAILURE       | Ethernet Out of Band Channel communications failure                                                       | 6              |
|           | Supervisor Hardware and CISCO_TAC | MGMT_PORT_FAILURE  | Hardware failure of management Ethernet port.                                                             | 5              |
|           | License                           | LICENSE_VIOLATION  | Feature in use that is not licensed (Release 1.3.x), and will be turned off after grace period expiration | 6              |
| Inventory | Inventory and CISCO_TAC           | COLD_BOOT          | Switch is powered up and reset to a cold boot sequence.                                                   | 2              |
|           |                                   | HARDWARE_INSERTION | New piece of hardware inserted into the chassis.                                                          | 2              |
|           |                                   | HARDWARE_REMOVAL   | Hardware removed from the chassis.                                                                        | 2              |
| Test      | Test and CISCO_TAC                | TEST               | User generated test.                                                                                      | 2              |

**Table 23-3 Event Categories and Command Outputs**

| Event Category            | Description                                                                                                                                                                                        | Executed Commands                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| System                    | Events generated by failure of a software system that is critical to unit operation.                                                                                                               | <b>show tech-support</b><br><b>show system redundancy status</b> |
| Environmental             | Events related to power, fan, and environment sensing elements such as temperature alarms.                                                                                                         | <b>show module</b><br><b>show environment</b>                    |
| Switching module hardware | Events related to standard or intelligent switching modules.                                                                                                                                       | <b>show tech-support</b>                                         |
| Supervisor hardware       | Events related to supervisor modules.                                                                                                                                                              | <b>show tech-support</b>                                         |
| Inventory                 | Inventory status is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This is considered a noncritical event, and the information is used for status and entitlement. | <b>show version</b>                                              |
| Test                      | User generated test message.                                                                                                                                                                       | <b>show version</b>                                              |

# Call Home Message Severity Levels

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family. Severity levels are preassigned per event type.

**Note**

Call Home severity levels are not the same as system message logging severity levels (see [Chapter 26, “Configuring System Message Logging”](#)).

Severity levels range from 0 to 9, with 9 having the highest urgency. Each severity level has keywords as listed in [Table 23-4](#).

**Table 23-4 Severity Levels**

| Severity Level | Keyword      | Description                                                                          |
|----------------|--------------|--------------------------------------------------------------------------------------|
| 9              | Catastrophic | Network wide catastrophic failure.                                                   |
| 8              | Disaster     | Significant network impact.                                                          |
| 7              | Fatal        | System is unusable.                                                                  |
| 6              | Critical     | Critical conditions, immediate attention needed.                                     |
| 5              | Major        | Major conditions.                                                                    |
| 4              | Minor        | Minor conditions.                                                                    |
| 3              | Warning      | Warning conditions.                                                                  |
| 2              | Notification | Basic notification and informational messages. Possibly independently insignificant. |
| 1              | Normal       | Normal event signifying return to normal state.                                      |
| 0              | Debugging    | Debugging messages.                                                                  |

# Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- E-mail address of the contact person
- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

[Table 23-5](#) describes the short text formatting option for all message types.

**Table 23-5 Short Text Messages**

| Data Item               | Description                                        |
|-------------------------|----------------------------------------------------|
| Device identification   | Configured device name                             |
| Date/time stamp         | Time stamp of the triggering event                 |
| Error isolation message | Plain English description of triggering event      |
| Alarm urgency level     | Error level such as that applied to syslog message |

[Table 23-6](#), [Table 23-7](#), and [Table 23-8](#) display the information contained in plain text and XML messages.

**Table 23-6 Reactive Event Message Format**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                               | XML Tag<br>(XML only) |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Time stamp                        | Date and time stamp of event in ISO time notation:<br><i>YYYY-MM-DDTHH:MM:SS</i> .<br><br><b>Note</b> The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time. | /mml/header/time      |
| Message name                      | Name of message. Specific event names are listed in the <a href="#">“Event Triggers”</a> section on page 23-13.                                                                                                                                   | /mml/header/name      |
| Message type                      | Specifically “Call Home”.                                                                                                                                                                                                                         | /mml/header/type      |
| Message group                     | Specifically “reactive”.                                                                                                                                                                                                                          | /mml/header/group     |
| Severity level                    | Severity level of message (see <a href="#">Table 23-4</a> ).                                                                                                                                                                                      | /mml/header/level     |
| Source ID                         | Product type for routing.                                                                                                                                                                                                                         | /mml/header/source    |

**Table 23-6 Reactive Event Message Format (continued)**

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>XML Tag<br/>(XML only)</b>    |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Device ID                                 | <p>Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p> | /mml/ header/deviceId            |
| Customer ID                               | Optional user-configurable field used for contract info or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                          | /mml/ header/customerID          |
| Contract ID                               | Optional user-configurable field used for contract info or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                          | /mml/ header /contractId         |
| Site ID                                   | Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.                                                                                                                                                                                                                                                                                                                                                                              | /mml/ header/siteId              |
| Server ID                                 | <p>If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch.</p> <p>Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p>                           | /mml/header/serverId             |
| Message description                       | Short text describing the error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | /mml/body/msgDesc                |
| Device name                               | Node that experienced the event. This is the host name of the device.                                                                                                                                                                                                                                                                                                                                                                                                                                | /mml/body/sysName                |
| Contact name                              | Name of person to contact for issues associated with the node experiencing the event.                                                                                                                                                                                                                                                                                                                                                                                                                | /mml/body/sysContact             |
| Contact e-mail                            | E-mail address of person identified as contact for this unit.                                                                                                                                                                                                                                                                                                                                                                                                                                        | /mml/body/sysContactEmail        |
| Contact phone number                      | Phone number of the person identified as the contact for this unit.                                                                                                                                                                                                                                                                                                                                                                                                                                  | /mml/body/sysContactPhone Number |
| Street address                            | Optional field containing street address for RMA part shipments associated with this unit.                                                                                                                                                                                                                                                                                                                                                                                                           | /mml/body/sysStreetAddress       |
| Model name                                | Model name of the switch. This is the specific model as part of a product family name.                                                                                                                                                                                                                                                                                                                                                                                                               | /mml/body/chassis/name           |
| Serial number                             | Chassis serial number of the unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | /mml/body/chassis/serialNo       |
| Chassis part number                       | Top assembly number of the chassis.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | /mml/body/chassis/partNo         |
| Chassis hardware version                  | Hardware version of chassis.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | /mml/body/chassis/hwVersion      |

**Table 23-6 Reactive Event Message Format (continued)**

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                  | <b>XML Tag<br/>(XML only)</b>      |
|-------------------------------------------|------------------------------------------------------------------------------|------------------------------------|
| Supervisor module software version        | Top level software version.                                                  | /mml/body/chassis/swVersion        |
| Affected FRU name                         | Name of the affected FRU generating the event message.                       | /mml/body/fru/name                 |
| Affected FRU serial number                | Serial number of affected FRU.                                               | /mml/body/fru/serialNo             |
| Affected FRU part number                  | Part number of affected FRU.                                                 | /mml/body/fru/partNo               |
| FRU slot                                  | Slot number of FRU generating the event message.                             | /mml/body/fru/slot                 |
| FRU hardware version                      | Hardware version of affected FRU.                                            | /mml/body/fru/hwVersion            |
| FRU software version                      | Software version(s) running on affected FRU.                                 | /mml/body/fru/swVersion            |
| Command output name                       | Exact command that was run. For example, <b>show running-config</b> command. | /mml/attachments/attachment/name   |
| Attachment type                           | Specifically command output.                                                 | /mml/attachments/attachment/type   |
| MIME type                                 | Normally text or plain or encoding type.                                     | /mml/attachments/attachment/mime   |
| Command output text                       | Output of command automatically executed (see <a href="#">Table 23-3</a> ).  | /mml/attachments/attachment/atdata |

**Table 23-7 Inventory Event Message Format**

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                                                                                                                                                                                    | <b>XML Tag<br/>(XML only)</b> |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| Time stamp                                | Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> .<br><br><b>Note</b> The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time. | /mml/header/time              |
| Message name                              | Name of message. Specifically “Inventory Update” Specific event names are listed in the <a href="#">“Event Triggers”</a> section on page 23-13.                                                                                                | /mml/header/name              |
| Message type                              | Specifically “Inventory Update”.                                                                                                                                                                                                               | /mml/header/type              |
| Message group                             | Specifically “proactive”.                                                                                                                                                                                                                      | /mml/header/group             |
| Severity level                            | Severity level of inventory event is level 2 (see <a href="#">Table 23-4</a> ).                                                                                                                                                                | /mml/header/level             |
| Source ID                                 | Product type for routing at Cisco. Specifically “MDS 9000”                                                                                                                                                                                     | /mml/header/source            |

**Table 23-7 Inventory Event Message Format (continued)**

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>XML Tag<br/>(XML only)</b>    |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Device ID                                 | <p>Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p> | /mml/ header /deviceId           |
| Customer ID                               | Optional user-configurable field used for contact info or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                                       | /mml/ header /customerID         |
| Contract ID                               | Optional user-configurable field used for contact info or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                                       | /mml/ header /contractId         |
| Site ID                                   | Optional user-configurable field, can be used for Cisco-supplied site ID or other data meaningful to alternate support service.                                                                                                                                                                                                                                                                                                                                                                                  | /mml/ header /siteId             |
| Server ID                                 | <p>If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch.</p> <p>Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p>                           | /mml/header/serverId             |
| Message description                       | Short text describing the error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | /mml/body/msgDesc                |
| Device name                               | Node that experienced the event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | /mml/body/sysName                |
| Contact name                              | Name of person to contact for issues associated with the node experiencing the event.                                                                                                                                                                                                                                                                                                                                                                                                                            | /mml/body/sysContact             |
| Contact e-mail                            | E-mail address of person identified as contact for this unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                    | /mml/body/sysContactEmail        |
| Contact phone number                      | Phone number of the person identified as the contact for this unit.                                                                                                                                                                                                                                                                                                                                                                                                                                              | /mml/body/sysContactPhone Number |
| Street address                            | Optional field containing street address for RMA part shipments associated with this unit.                                                                                                                                                                                                                                                                                                                                                                                                                       | /mml/body/sysStreetAddress       |
| Model name                                | Model name of the unit. This is the specific model as part of a product family name.                                                                                                                                                                                                                                                                                                                                                                                                                             | /mml/body/chassis/name           |
| Serial number                             | Chassis serial number of the unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | /mml/body/chassis/serialNo       |
| Chassis part number                       | Top assembly number of the chassis.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | /mml/body/chassis/partNo         |
| Chassis hardware version                  | Hardware version of chassis.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | /mml/body/chassis/hwVersion      |

**Table 23-7 Inventory Event Message Format (continued)**

| Data Item<br>(Plain text and XML)  | Description<br>(Plain text and XML)                                                                                            | XML Tag<br>(XML only)              |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| Supervisor module software version | Top level software version.                                                                                                    | /mml/body/chassis/swVersion        |
| FRU name                           | Name of the affected FRU generating the event message.                                                                         | /mml/body/fru/name                 |
| FRU s/n                            | Serial number of FRU.                                                                                                          | /mml/body/fru/serialNo             |
| FRU part number                    | Part number of FRU.                                                                                                            | /mml/body/fru/partNo               |
| FRU slot                           | Slot number of FRU.                                                                                                            | /mml/body/fru/slot                 |
| FRU hardware version               | Hardware version of FRU.                                                                                                       | /mml/body/fru/hwVersion            |
| FRU software version               | Software version(s) running on FRU.                                                                                            | /mml/body/fru/swVersion            |
| Command output name                | Exact command that was run. For example, the <b>show running-config</b> command.                                               | /mml/attachments/attachment/name   |
| Attachment type                    | Specifically command output.                                                                                                   | /mml/attachments/attachment/type   |
| MIME type                          | Normally text or plain or encoding type.                                                                                       | /mml/attachments/attachment/mime   |
| Command output text                | Output of command automatically executed after event categories (see <a href="#">“Event Triggers” section on page 23-13</a> ). | /mml/attachments/attachment/atdata |

**Table 23-8 User-Generated Test Message Format**

| Data Item<br>(Plain text and XML) | Description<br>(Plain text and XML)                                                                                                                                                                                                            | XML Tag<br>(XML only) |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Time stamp                        | Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> .<br><br><b>Note</b> The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time. | /mml/header/time      |
| Message name                      | Name of message. Specifically test message for test type message. Specific event names listed in the <a href="#">“Event Triggers” section on page 23-13</a> ).                                                                                 | /mml/header/name      |
| Message type                      | Specifically “Test Call Home”.                                                                                                                                                                                                                 | /mml/header/type      |
| Message group                     | This field should be ignored by the receiving Call Home processing application, but may be populated with either “proactive” or “reactive”.                                                                                                    | /mml/header/group     |
| Severity level                    | Severity level of message, test Call Home message (see <a href="#">Table 23-4</a> ).                                                                                                                                                           | /mml/header/level     |
| Source ID                         | Product type for routing.                                                                                                                                                                                                                      | /mml/header/source    |



**Table 23-8 User-Generated Test Message Format (continued)**

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>XML Tag<br/>(XML only)</b>          |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| Device ID                                 | <p>Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p> | /mml/ header /deviceId                 |
| Customer ID                               | Optional user-configurable field used for contract info or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                                      | /mml/ header /customerId               |
| Contract ID                               | Optional user-configurable field used for contract info or other ID by any support service.                                                                                                                                                                                                                                                                                                                                                                                                                      | /mml/ header /contractId               |
| Site ID                                   | Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.                                                                                                                                                                                                                                                                                                                                                                                          | /mml/ header /siteId                   |
| Server ID                                 | <p>If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch.</p> <p>Format: type@Sid@serial, where</p> <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> <p>Example: “DS-C9000@C@12345678</p>                           | /mml/header/serverId                   |
| Message description                       | Short text describing the error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | /mml/body/msgDesc                      |
| Device name                               | Switch that experienced the event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | /mml/body/sysName                      |
| Contact name                              | Name of person to contact for issues associated with the node experiencing the event.                                                                                                                                                                                                                                                                                                                                                                                                                            | /mml/body/sysContact                   |
| Contact Email                             | E-mail address of person identified as contact for this unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                    | /mml/body/sysContactEmail              |
| Contact phone number                      | Phone number of the person identified as the contact for this unit.                                                                                                                                                                                                                                                                                                                                                                                                                                              | /mml/body/sysContactPhone<br>Number    |
| Street address                            | Optional field containing street address for RMA part shipments associated with this unit.                                                                                                                                                                                                                                                                                                                                                                                                                       | /mml/body/sysStreetAddress             |
| Model name                                | Model name of the switch. This is the specific model as part of a product family name.                                                                                                                                                                                                                                                                                                                                                                                                                           | /mml/body/chassis/name                 |
| Serial number                             | Chassis serial number of the unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | /mml/body/chassis/serialNo             |
| Chassis part number                       | Top assembly number of the chassis. For example, 800-xxx-xxxx.                                                                                                                                                                                                                                                                                                                                                                                                                                                   | /mml/body/chassis/partNo               |
| Command output text                       | Output of command automatically executed after event categories listed in <a href="#">Table 23-3</a> .                                                                                                                                                                                                                                                                                                                                                                                                           | /mml/attachments/attachmen<br>t/atdata |

**Table 23-8** *User-Generated Test Message Format (continued)*

| <b>Data Item<br/>(Plain text and XML)</b> | <b>Description<br/>(Plain text and XML)</b>                                      | <b>XML Tag<br/>(XML only)</b>    |
|-------------------------------------------|----------------------------------------------------------------------------------|----------------------------------|
| MIME type                                 | Normally text or plain or encoding type.                                         | /mml/attachments/attachment/mime |
| Attachment type                           | Specifically command output.                                                     | /mml/attachments/attachment/type |
| Command output name                       | Exact command that was run. For example, the <b>show running-config</b> command. | /mml/attachments/attachment/name |



## Configuring Domain Parameters

---

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per VSAN basis.

This chapter includes the following sections:

- [About fcdomain Phases, page 24-2](#)
- [Restarting the Domain, page 24-3](#)
- [Configuring the Domain, page 24-4](#)
- [Setting Switch Priority, page 24-6](#)
- [Configuring Allowed Domain ID Lists, page 24-6](#)
- [Merging Stable Fabrics, page 24-7](#)
- [Assigning Contiguous Domains, page 24-7](#)
- [Disabling the fcdomain Feature, page 24-8](#)
- [Setting the Fabric Name, page 24-8](#)
- [Stopping Incoming RCFs, page 24-9](#)
- [Enabling Persistent FC IDs, page 24-9](#)
- [Configuring Persistent FC IDs Manually, page 24-10](#)
- [Purging Persistent FC IDs, page 24-13](#)
- [Displaying fcdomain Information, page 24-13](#)
- [Default Settings, page 24-16](#)



### Caution

Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.



### Tip

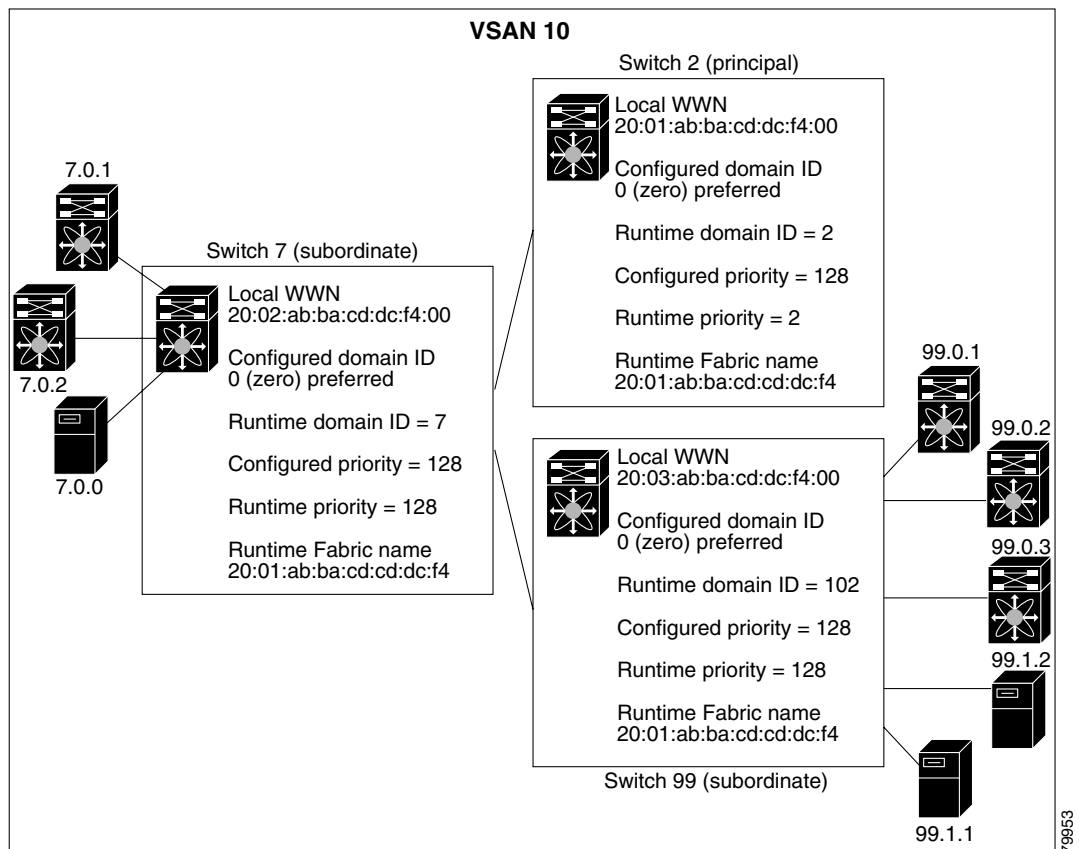
When you change the configuration, be sure to save the running configuration using the **copy running-config startup-config** command. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.

## About fcdomain Phases

This section describes each fcdomain phase (see [Figure 24-1](#)):

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- Fabric reconfiguration—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

**Figure 24-1 Sample fcdomain Configuration**



### Note

Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

# Restarting the Domain

The **fcdomain restart** command applies your changes to the runtime settings. Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric. If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric.

To restart the fabric disruptively or nondisruptively, follow these steps:

|        | Command                                                   | Purpose                                                      |
|--------|-----------------------------------------------------------|--------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                | Enters configuration mode.                                   |
| Step 2 | switch(config)# <b>fcdomain restart vsan 1</b>            | Forces the VSAN to reconfigure without traffic disruption.   |
|        | switch(config)# <b>fcdomain restart disruptive vsan 1</b> | Forces the VSAN to reconfigure with data traffic disruption. |

You can apply most of the configurations to their corresponding runtime values by using the **restart disruptive** option. Each of the following sections provide further details on how the **fcdomain** parameters are applied to the runtime values.

**Note**

A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID will change to take on the static domain ID after the next restart (see the [“Configuring the Domain” section on page 24-4](#)).

**Tip**

If a VSAN is in **interop** mode, you cannot restart the **fcdomain** for that VSAN disruptively.

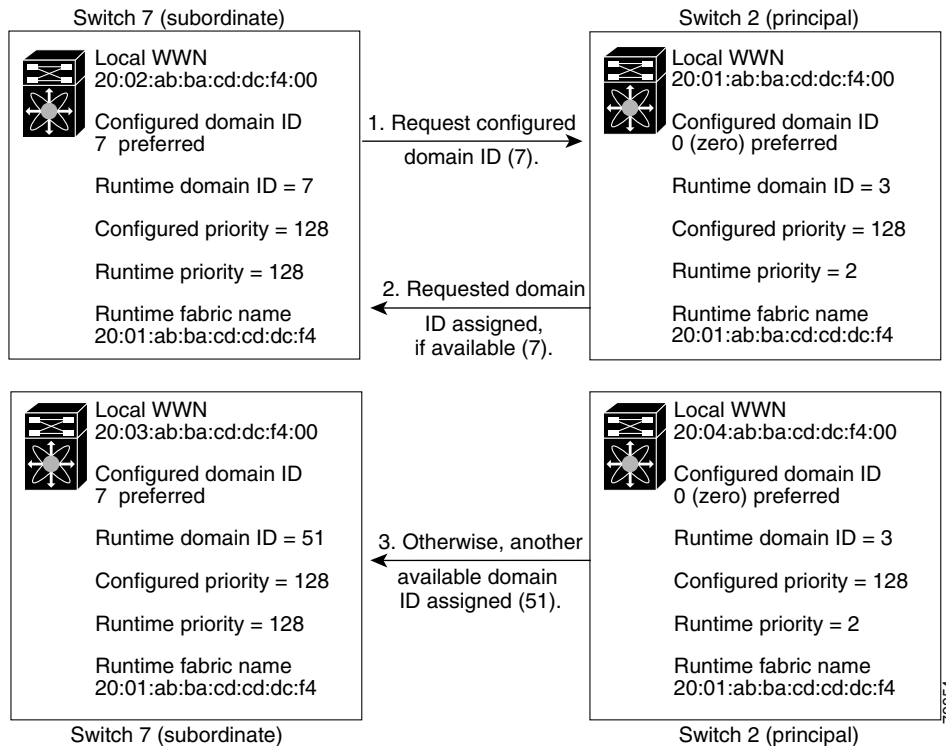
# Configuring the Domain

The configured domain ID can be **preferred** or **static**. By default, the configured domain is **0** and the configured type is **preferred**. If you do not configure a domain ID, the local switch sends a random ID in its request.

When a subordinate switch requests a domain, the following process takes place (see Figure 24-2):

1. The local switch sends a configured domain ID request to the principal switch.
2. The principal switch assigns the requested domain ID if available, otherwise, it assigns another available domain ID.

**Figure 24-2 Configuration Process Using the preferred Option**



The behavior for a subordinate switch changes based on the allowed domain ID lists, on the configured domain ID, and on the domain ID that the principal switch has assigned to the requesting switch:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the **preferred** and **static** options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
  - If the configured type is **static**, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
  - If the configured type is **preferred**, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

**Caution**

You must issue the **fcdomain restart** command if you want to apply the configured domain changes to the runtime domain.

To specify a **preferred** or a **static** domain ID, follow these steps:

|               | Command                                                      | Purpose                                                                                                                                                                      |
|---------------|--------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>config t</b><br>switch(config)#                   | Enters configuration mode.                                                                                                                                                   |
| <b>Step 2</b> | switch(config)# <b>fcdomain domain 3 preferred vsan 8</b>    | Configures the switch in VSAN 8 to request a preferred domain ID 3 and accepts any value assigned by the principal switch.                                                   |
|               | switch(config)# <b>no fcdomain domain 3 preferred vsan 8</b> | Resets the configured domain ID to 0 (default) in VSAN 8. The configured domain ID becomes 0 preferred.                                                                      |
| <b>Step 3</b> | switch(config)# <b>fcdomain domain 2 static vsan 237</b>     | Configures the switch in VSAN 237 to accept only a specific value and moves the local interfaces in VSAN 237 to an isolated state if the requested domain ID is not granted. |
|               | switch(config)# <b>no fcdomain domain 18 static vsan 237</b> | Resets the configured domain ID to factory defaults in VSAN 237. The configured domain ID becomes 0 preferred.                                                               |

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.

**Note**

The 0 (zero) value can be configured only if you use the **preferred** option.

While the **static** option can be applied to runtime after a disruptive or nondisruptive restart, the **preferred** option is applied to runtime only after a disruptive restart (see the [“Restarting the Domain” section on page 24-3](#)).

**Tip**

When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN will remain in the static state. You can change the static ID value but you cannot change it to the preferred option.

## Setting Switch Priority

By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

Any new switch cannot become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower WWN becomes the principal switch.

To configure the priority for the principal switch, follow these steps:

|        | Command                                                | Purpose                                                       |
|--------|--------------------------------------------------------|---------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#             | Enters configuration mode.                                    |
| Step 2 | switch(config)# <b>fcdomain priority 25 VSAN 99</b>    | Configures a priority of 25 for the local switch in VSAN 99.  |
|        | switch(config)# <b>no fcdomain priority 25 VSAN 99</b> | Reverts the priority to the factory default (128) in VSAN 99. |

The priority configuration is applied to runtime through a disruptive restart (see the [“Restarting the Domain”](#) section on page 24-3).

## Configuring Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch ensures that the domain requested by any switch in the fabric is specified in the allowed list.



### Tip

If you configure an allowed list on one switch in the fabric, we recommend you configure the same list in all other switches in the fabric to ensure consistency.

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.
- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally-configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already-configured domain ID lists must not be empty.

To configure the allowed domain ID list, follow these steps:

|        | Command                                                  | Purpose                                                                             |
|--------|----------------------------------------------------------|-------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#               | Enters configuration mode.                                                          |
| Step 2 | switch(config)# <b>fcdomain allowed 50-110 vsan 4</b>    | Configures the list to allow switches with the domain ID 50 through 110 in VSAN 4.  |
|        | switch(config)# <b>no fcdomain allowed 50-110 vsan 5</b> | Reverts to the factory default of allowing domain IDs from 1 through 239 in VSAN 5. |



## Merging Stable Fabrics

By default, the **auto-reconfigure** option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

- If the **auto-reconfigure** option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the **auto-reconfigure** option is disabled on either or both switches, the links between the two switches become isolated.

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), follow these steps:

|        | Command                                                  | Purpose                                                                                         |
|--------|----------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#               | Enters configuration mode.                                                                      |
| Step 2 | switch(config)# <b>fcdomain auto-reconfigure vsan 10</b> | Enables the automatic reconfiguration option in VSAN 10.                                        |
|        | switch(config)# <b>no fcdomain auto-reconfigure 69</b>   | Disables the automatic reconfiguration option and reverts it to the factory default in VSAN 69. |

The **auto-reconfigure** option takes immediate effect at runtime—you do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the **auto-reconfigure** option on both switches, the fabric continues to be isolated. However, if you enable the **auto-reconfigure** option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) occurs. A disruptive reconfiguration may affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and getting rid of the domain overlap.

## Assigning Contiguous Domains

By default, the **contiguous-allocation** option is disabled. When the subordinate switches request the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the **contiguous-allocation** option is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches.
- If the **contiguous-allocation** option is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switches.

To enable contiguous domains in a specific VSAN (or a range of VSANs), follow these steps:

|        | Command                                                            | Purpose                                                                                       |
|--------|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                         | Enters configuration mode.                                                                    |
| Step 2 | switch(config)# <b>fcdomain contiguous-allocation vsan 81-83</b>   | Enables the contiguous allocation option in VSAN 81 through 83.                               |
|        | switch(config)# <b>no fcdomain contiguous-allocation vsan 1030</b> | Disables the contiguous allocation option and reverts it to the factory default in VSAN 1030. |

The **contiguous-allocation** option takes immediate effect at runtime—you do not need to restart the fcdomain.

## Disabling the fcdomain Feature

By default, the fcdomain feature is enabled on each switch. You can disable the fcdomain feature by using the **no fcdomain** command. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric.

To disable fcdomains in a single VSAN or a range of VSANs, follow these steps:

|        | Command                                       | Purpose                                                    |
|--------|-----------------------------------------------|------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#    | Enters configuration mode.                                 |
| Step 2 | switch(config)# <b>no fcdomain vsan 7-200</b> | Disables the fcdomain configuration in VSAN 7 through 200. |
|        | switch(config)# <b>fcdomain vsan 2008</b>     | Enables the fcdomain configuration in VSAN 2008.           |

The fcdomain configuration is applied to runtime through a disruptive restart.

## Setting the Fabric Name

By default the configured fabric name is 20:01:00:05:30:00:28:df.

- When the fcdomain feature is disabled, the runtime fabric name is the same as the configured fabric name.
- When the fcdomain feature is enabled, the runtime fabric name is the same as the principal switch's WWN.

To set the fabric name value for a disabled fcdomain, follow these steps:

|        | Command                                                                        | Purpose                                                                                      |
|--------|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                     | Enters configuration mode.                                                                   |
| Step 2 | switch(config)# <b>fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 3</b>       | Assigns the configured fabric name value in VSAN 3.                                          |
|        | switch(config)# <b>no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 3010</b> | Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010. |

The fabric name is applied to runtime through a disruptive restart when the fcdomain is configured as disabled (see the [“Restarting the Domain”](#) section on page 24-3).

## Stopping Incoming RCFs

The **rcf-reject** option is configured on a per-interface, per-VSAN basis. By default, the **rcf-reject** option is disabled (that is, RCF request frames are not automatically rejected).

To stop incoming RCF request frames, follow these steps:

|        | Command                                                 | Purpose                                                       |
|--------|---------------------------------------------------------|---------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#              | Enters configuration mode.                                    |
| Step 2 | switch(config)# <b>int fc1/1</b><br>switch(config-if)#  | Configures the specified interface.                           |
| Step 3 | switch(config-if)# <b>fcdomain rcf-reject vsan 1</b>    | Enables the RCF filter on the specified interface in VSAN 1.  |
|        | switch(config-if)# <b>no fcdomain rcf-reject vsan 1</b> | Disables the RCF filter on the specified interface in VSAN 1. |

The **rcf-reject** option takes immediate effect to runtime through a disruptive restart (see the “[Restarting the Domain](#)” section on page 24-3).

## Enabling Persistent FC IDs

Persistent FC IDs are disabled by default. You can enable this option globally or for each VSAN. If you choose to enable it globally, you can do so at any time using the initial setup routine or the setup command (see the “[Initial Setup Routine](#)” section on page 4-2). When you enable this option globally, the switch remains in this state until you change the global configuration.



### Note

If you enable this option during the initial switch setup, this option will be automatically enabled in all configured VSANs. If you enable this option at a later stage, this option will be automatically enabled in all VSANs configured after that stage. VSANs configured before that stage will remain unchanged.

When a N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned a FC ID. By default, the persistent FC ID feature is disabled. If this feature is disabled, the following consequences apply:

- A N or NL port logs into a Cisco MDS 9000 Family switch, the WWN of the requesting N or NL port and the assigned FC ID, are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN, on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted, and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- The switch connection behavior differs between N ports and NL ports:
  - N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).
  - NL ports receive the same FC IDs only if connected back to the same port on the switch to which it was originally connected.

The assigned FC IDs in a `fcdomain` can be enabled to remain persistent even after a reboot. This ensures that an attached N port receives the same FC IDs after a reboot. If you enable this feature, the following consequences apply:

- The currently *in-use* FC IDs in the `fcdomain` are saved across reboots.
- The `fcdomain` automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.



**Note** If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts.

To enable the persistent FC ID feature, follow these steps:

|        | Command                                                                                                   | Purpose                                            |
|--------|-----------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Step 1 | <code>switch# config t</code>                                                                             | Enters configuration mode.                         |
| Step 2 | <code>switch(config)# fcdomain fcid persistent vsan 1000</code><br>FCID(s) persistent feature is enabled. | Activates persistency of FC IDs in VSAN 1000.      |
|        | <code>switch(config-if)# no fcdomain fcid persistent vsan 20</code>                                       | Disables the FC ID persistency feature in VSAN 20. |



**Note** Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

## Configuring Persistent FC IDs Manually

Once the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis. Follow these requirements to manually configure a persistent FC ID:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.
- Ensure that the require VSAN is an active VSAN—persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.
- Do not replace an FC ID that is already configured in another WWN. If you want to use a previously-configured WWN, first delete the configured WWN before proceeding with this procedure.

To configure persistent FC IDs, follow these steps:

|        | Command                                                                                    | Purpose                                                                                                                                                                                                                      |
|--------|--------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#                                                 | Enters configuration mode.                                                                                                                                                                                                   |
| Step 2 | switch(config)# <b>fcdomain fcid database</b>                                              | Activates persistency of FC IDs in the specified VSAN.                                                                                                                                                                       |
| Step 3 | switch(config-fcid-db)# <b>vsan 1000 wwn 33:e8:00:05:30:00:16:df fcid 0x070128</b>         | Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in VSAN 1000.                                                                                                                                      |
|        | switch(config-fcid-db)# <b>vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070123 dynamic</b> | Configures a device WWN (11:22:11:22:33:44:33:44) with the FC ID 0x070123 in VSAN 1000 in dynamic mode.                                                                                                                      |
|        | switch(config-fcid-db)# <b>vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070100 area</b>    | Configures a device WWN (11:22:11:22:33:44:33:44) with the FC IDs 0x070100 through 0x0701FF in VSAN 1000.<br><br><b>Note</b> To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID. |

**Note**

You cannot configure persistent FC IDs in FICON-enabled VSANs.

## Configuring Unique Area FC IDs for Some HBAs

**Note**

This section does not apply if either the HBA port or the storage port is connected to different switches.

Some HBA ports require a different area ID than storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Switches in the Cisco MDS 9000 Family facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port. The procedure in this example uses a switch domain of **111(6f hex)**. The HBA port connects to **interface fc1/9** and the storage port connects to **interface fc 1/10** in the same switch.

To configure a different area ID for the HBA port, follow these steps:

**Step 1** Obtain the Port WWN (Port Name field) ID of the HBA using the **show flogi database** command).

```
switch# show flogi database
```

```

INTERFACE VSAN FCID PORT NAME NODE NAME

fc1/9 3 0x6f7703 50:05:08:b2:00:71:c8:c2 50:05:08:b2:00:71:c8:c0
fc1/10 3 0x6f7704 50:06:0e:80:03:29:61:0f 50:06:0e:80:03:29:61:0f

```



**Note** Both FC IDs in this setup have the same area 77 assignment.

**Step 2** Shut down the HBA interface in the MDS switch.

```
switch# conf t
switch(config)# interface fc1/9
switch(config-if)# shutdown
switch(config-if)# end
switch#
```

**Step 3** Verify that the FC ID feature is enabled using the show fcdomain vsan command.

```
switch# show fcdomain vsan 1
...
Local switch configuration information:
 State: Enabled
 FCID persistence: Disabled
```

If this feature is disabled, continue with this procedure to enable the FC ID persistence.

If this feature is already enabled, skip to [Step 5](#).

**Step 4** Enable the FC ID persistence feature in the MDS switch.

```
switch# conf t
switch(config)# fcdomain fcid persistent vsan 1
switch(config)# end
switch#
```

**Step 5** Assign a new FC ID with a different area allocation. In this example, we replace 77 with *ee*.

```
switch# conf t
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan 3 wwn 50:05:08:b2:00:71:c8:c2 fcid 0x6fee00
```

**Step 6** Enable the HBA interface in the MDS switch.

```
switch# conf t
switch(config)# interface fc1/9
switch(config-if)# no shutdown
switch(config-if)# end
switch#
```

**Step 7** Verify the pWWN ID of the HBA using the show flogi database command.

```
switch# show flogi database
```

| INTERFACE | VSAN | FCID     | PORT NAME               | NODE NAME               |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/9     | 3    | 0x6fee00 | 50:05:08:b2:00:71:c8:c2 | 50:05:08:b2:00:71:c8:c0 |
| fc1/10    | 3    | 0x6f7704 | 50:06:0e:80:03:29:61:0f | 50:06:0e:80:03:29:61:0f |



**Note** Both FC IDs now have different area assignments.

## Purging Persistent FC IDs

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. [Table 24-1](#) identifies the FC ID entries that are deleted by the **purge fcdomain** command.

**Table 24-1 Purged FC IDs**

| Persistent FC ID state | Persistent Usage State | Action      |
|------------------------|------------------------|-------------|
| static                 | in use                 | Not deleted |
| static                 | not in use             | Not deleted |
| dynamic                | in use                 | Not deleted |
| dynamic                | not in use             | deleted     |

Dynamic, not in use, FC IDs can be removed using the **purge fcdomain** command (see [Table 24-1](#)).

To purge persistent FC IDs, follow this step:

|        | Command                                     | Purpose                                                   |
|--------|---------------------------------------------|-----------------------------------------------------------|
| Step 1 | switch# <b>purge fcdomain fcid vsan 4</b>   | Purges all dynamic and unused FC IDs in VSAN 4            |
|        | switch# <b>purge fcdomain fcid vsan 3-5</b> | Purges all dynamic and unused FC IDs in VSAN 3, 4, and 5. |

## Displaying fcdomain Information

The **show fcdomain** commands display global information about the fcdomain configurations. See [Example 24-1](#).



### Note

In [Example 24-1](#), the fcdomain feature is disabled. Consequently, the runtime fabric name is the same as the configured fabric name.

### Example 24-1 Displays the Global fcdomain Information

```
switch# show fcdomain vsan 2
The local switch is the Principal Switch.

Local switch run time information:
 State: Stable
 Local switch WWN: 20:01:00:0b:46:79:ef:41
 Running fabric name: 20:01:00:0b:46:79:ef:41
 Running priority: 128
 Current domain ID: 0xed(237)

Local switch configuration information:
 State: Enabled
 FCID persistence: Disabled
 Auto-reconfiguration: Disabled
 Contiguous-allocation: Disabled
 Configured fabric name: 20:01:00:05:30:00:28:df
 Configured priority: 128
 Configured domain ID: 0x00(0) (preferred)
```

```
Principal switch run time information:
 Running priority: 128
```

```
No interfaces available.
```

Use **show fcdomain domain-list** command to display the list of domain IDs of all switches belonging to a specified VSAN. This list provides the WWN of the switches owning each domain ID. See [Example 24-2](#).

#### Example 24-2 Displays the fcdomain Lists

```
switch# show fcdomain domain-list vsan 1

Number of domains: 1
Domain ID WWN

0x16 (22) 20:01:00:05:30:00:16:df [Local] [Principal]
```

Use **show** command to display the list of allowed domain IDs configured on this switch. See [Example 24-3](#).

#### Example 24-3 Displays the Allowed Domain ID Lists

```
switch# show fcdomain allowed vsan 1
Assigned or unallowed domain IDs: 1-96,100,111-239.
[Interoperability Mode 1] allowed domain IDs: 97-127.
[User] configured allowed domain IDs: 50-110.
```



#### Tip

Ensure that the requested domain ID passes the SAN-OS software checks, if **interop 1** mode is required in this switch.

Use the **show fcdomain fcid persistent** command to display all existing, persistent FC IDs for a specified VSAN. You can also specify the **unused** option to view only persistent FC IDs that are still not in use. See Examples [24-4](#) and [24-5](#).

#### Example 24-4 Displays Persistent FC IDs in a Specified VSAN

```
switch# show fcdomain fcid persistent vsan 1000
Total entries 2.

Persistent FCIDs table contents:
VSAN WWN FCID Mask Used Assignment

1000 11:11:22:22:11:11:12:23 0x700101 SINGLE FCID NO STATIC
1000 44:44:33:33:22:22:11:11 0x701000 ENTIRE AREA NO DYNAMIC
```

#### Example 24-5 Displays All Persistent FC IDs in the fcdomain

```
switch# show fcdomain fcid persistent
Total entries 2.

Persistent FCIDs table contents:
VSAN WWN FCID Mask Used Assignment

```



|      |                         |          |             |     |         |
|------|-------------------------|----------|-------------|-----|---------|
| 1000 | 11:11:22:22:11:11:22:22 | 0x700501 | SINGLE FCID | NO  | STATIC  |
| 1003 | 44:44:33:33:22:22:11:11 | 0x781000 | ENTIRE AREA | YES | DYNAMIC |

Use the **show fcdomain statistics** command to display frame and other fcdomain statistics, for a specified VSAN or PortChannel. See [Example 24-6](#) and [Example 24-7](#).

**Example 24-6 Displays fcdomain Statistics for a Specified VSAN**

```
switch# show fcdomain statistics vsan 1
VSAN Statistics
 Number of Principal Switch Selections: 5
 Number of times Local Switch was Principal: 0
 Number of 'Build Fabric's: 3
 Number of 'Fabric Reconfigurations': 0
```

**Example 24-7 Displays fcdomain Statistics for a Specified PortChannel**

```
switch# show fcdomain statistics interface port-channel 10 vsan 1
Interface Statistics:
 Transmitted Received

 EFPs 13 9
 DIAs 7 7
 RDIs 0 0
 ACCs 21 25
 RJTs 1 1
 BFs 2 2
 RCFs 4 4
 Error 0 0
 Total 48 48
Total Retries: 0
Total Frames: 96

```

Use the **show fcdomain address-allocation** command to display FC ID allocation statistics including a list of assigned and free FC IDs. See [Example 24-8](#).

**Example 24-8 Displays FC ID Information**

```
switch# show fcdomain address-allocation vsan 1
Free FCIDs: 0x020000 to 0x02fdff
 0x02ff00 to 0x02fffe

Assigned FCIDs: 0x02fe00 to 0x02feff
 0x02ffff

Reserved FCIDs: 0x020100 to 0x02f0ff
 0x02fe00 to 0x02feff
 0x02ffff

Number free FCIDs: 65279
Number assigned FCIDs: 257
Number reserved FCIDs: 61697
```

Use the **show fcdomain address-allocation cache** command to display the valid address-allocation cache. The cache is used by the principal switch to reassign the FC IDs for a device (disk or host) that exited and reentered the fabric. In the cache content, VSAN refers to the VSAN that contains the device, WWN refers to the device that owned the FC IDs, and mask refers to a single or entire area of FC IDs. See [Example 24-9](#).

**Example 24-9 Displays Address Allocation Information**

```

switch# show fcdomain address-allocation cache
Cache content:
line# VSAN WWN FCID mask

1. 12 21:00:00:e0:8b:08:a2:21 0xef0400 ENTIRE AREA
2. 6 50:06:04:82:c3:a1:2f:5c 0xef0002 SINGLE FCID
3. 8 20:4e:00:05:30:00:24:5e 0xef0300 ENTIRE AREA
4. 8 50:06:04:82:c3:a1:2f:52 0xef0001 SINGLE FCID

```

## Default Settings

Table 24-2 lists the default settings for all fcdomain parameters.

**Table 24-2 Default fcdomain Parameters**

| Parameters                   | Default                           |
|------------------------------|-----------------------------------|
| fcdomain feature             | Enabled.                          |
| Configured domain ID         | 0 (zero).                         |
| Configured domain option     | Preferred.                        |
| auto-reconfigure option      | Disabled.                         |
| contiguous-allocation option | Disabled.                         |
| Priority                     | 128.                              |
| Allowed list                 | 1 to 239.                         |
| Fabric-name                  | 20:01:00:05:30:00:28:df.          |
| rcf-reject                   | Disabled.                         |
| Persistent FC ID             | Disabled (globally configurable). |



## Configuring Traffic Management

---

Fibre Channel Congestion Control (FCC) is a Cisco proprietary flow control mechanism that alleviates congestion on Fibre Channel networks.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.
- Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

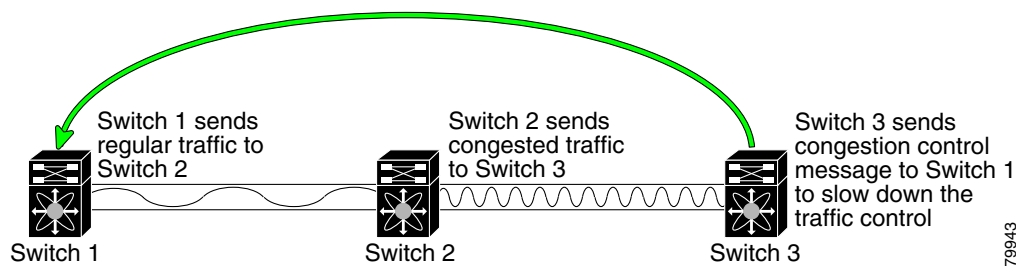
This chapter provides details on the QoS and FCC features provided in all switches. It includes the following sections:

- [FCC, page 25-2](#)
- [QoS, page 25-4](#)
- [Control Traffic, page 25-4](#)
- [Data Traffic, page 25-5](#)
- [Ingress Port Rate Limiting, page 25-12](#)
- [Default Settings, page 25-12](#)

# FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic (see [Figure 25-1](#)).

**Figure 25-1 FCC Mechanisms**



Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

## FCC Process

When a node in the network detects a congestion for an output port, it generates an edge quench message. These frames are identified by the Fibre Channel destination ID (DID) and the source ID. A switch from other vendors simply forwards these frames.

Any receiving switch in the Cisco MDS 9000 Family handles frames in one of these ways:

- It forwards the frame.
- It limits the rate of the frame flow in the congested port.

The behavior of the flow control mechanism differs based on the Fibre Channel DID:

- If the Fibre Channel DID is directly connected to one of the switch ports, the input rate limit is applied to that port.
- If the destination of the edge quest frame is a Cisco domain or the next hop is a Cisco MDS 9000 Family switch, the frame is forwarded.
- If neither of these mechanisms is true, then the frame is processed in the port going towards the FC DID.

All switches (including the edge switch) along the congested path process path quest frames. However, only the edge switch processes edge quest frames.

## Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled for the entire switch.

**Tip**

If you enable FCC, be sure to enable it in all switches in the fabric.

To enable or disable the FCC feature, follow these steps:

|        | Command                       | Purpose                                |
|--------|-------------------------------|----------------------------------------|
| Step 1 | switch# <b>config t</b>       | Enters configuration mode.             |
| Step 2 | switch(config)# <b>fcc</b>    | Enables FCC in this switch.            |
|        | switch(config)# <b>no fcc</b> | Disables FCC in this switch (default). |

## Assigning FCC Priority

To assign FCC priority, follow these steps:

|        | Command                               | Purpose                                                                                                            |
|--------|---------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b>               | Enters configuration mode.                                                                                         |
| Step 2 | switch(config)# <b>fcc priority 2</b> | Defines the FCC priority threshold to have a priority of 2—0 is the lowest priority and 7 is the highest priority. |

## Displaying FCC

Use the show fcc commands to view FCC settings (see [Example 25-1](#)).

**Example 25-1 Displays Configured FCC Information**

```
switch# show fcc
fcc is disabled
fcc is applied to frames with priority up to 4
```

# QoS

QoS implementation in the Cisco MDS 9000 Family follows the Differentiated Services (DiffServ) model. The DiffServ standard is defined in RFCs 2474 and 2475.

All switches support the following types of traffic:

- Control Traffic, [page 25-4](#)
- Data Traffic, [page 25-5](#)

## Control Traffic

The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:

- Internally generated time-critical control traffic (mostly Class F frames).
- Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor’s switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

## Disabling Control Traffic

By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.



Tip

We do not recommended disabling this feature as all critical control traffic will automatically be assigned the lowest priority once you issue this command. You can view the current state of the QoS configuration for critical control traffic using the **show qos statistics** command.

To disable the high priority assignment for control traffic, follow these steps:

|        | Command                                          | Purpose                                   |
|--------|--------------------------------------------------|-------------------------------------------|
| Step 1 | switch# <b>config t</b>                          | Enters configuration mode.                |
| Step 2 | switch(config)# <b>no qos control priority 0</b> | Enables the control traffic QoS feature.  |
|        | switch(config)# <b>qos control priority 0</b>    | Disables the control traffic QoS feature. |

## Displaying Control Traffic Information

The **show qos** command displays the current QoS settings along with a the number of frames marked high priority. The count is only for debugging purposes and cannot be configured (see [Example 25-2](#)).

### Example 25-2 Displays Current QoS Settings

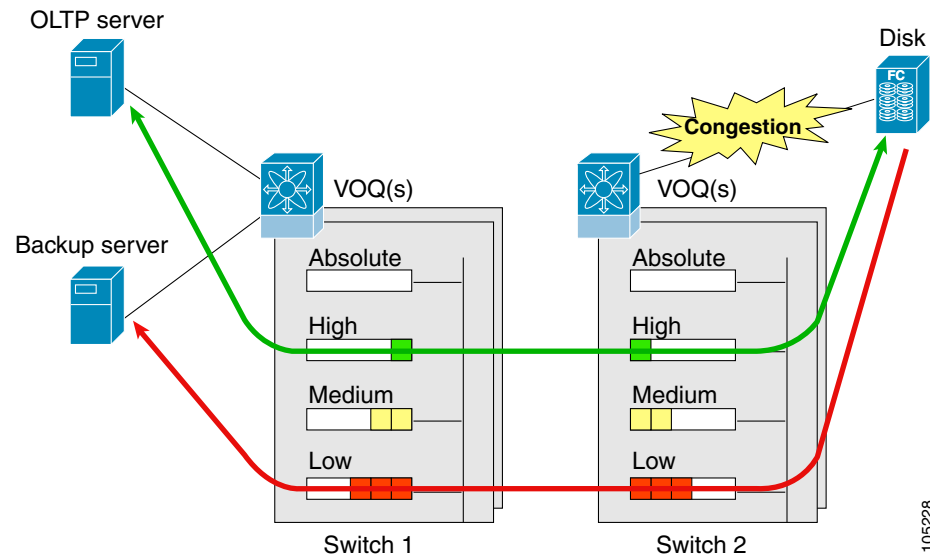
```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 15767
Number of highest-priority FC frames transmitted = 8224
Current priority of FC control frames = 0 (0 = lowest; 7 = highest)
```

# Data Traffic

Transaction processing, a low volume, latency sensitive application, requires quick access to requested information. Backup processing requires high bandwidth but is not sensitive to latency. In a network that does not support service differentiation, all traffic is treated identically—they experience similar latency and get similar bandwidths. The QoS feature in all switches in the Cisco MDS 9000 Family provides these guarantees from SAN-OS Release 1.3(x).

Prior versions of the SAN-OS software only differentiated traffic priority based on control traffic. SAN-OS Release 1.3(x) enables you to take full advantage of the QoS capabilities. Data traffic can now be prioritized in four distinct levels of service differentiation: low, medium, high, or absolute priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications like data warehousing (see [Figure 25-1](#)).

**Figure 25-2 Prioritizing Data Traffic**



In [Figure 25-1](#), the OLTP traffic arriving at Switch 1 is marked with a **High** priority level of through classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a **Low** priority level. The traffic is sent to the corresponding priority queue within a Virtual Output Queue (VOQ).

A Deficit Weighted Round Robin (DWRR) scheduler configured in the first switch ensures that high priority traffic is treated better than low priority traffic. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

If the ISL is congested when the OLTP server sends a request, the request is queued in the high priority queue and is serviced almost immediately as the high priority queue is not congested. The scheduler assigns it priority over the backup traffic in the low priority queue.



## Note

When the high priority queue does not have traffic flowing through, the low priority queue uses all the bandwidth and is not restricted to the configured value.

A similar occurrence in Switch 2 sends a response to the transaction request. The round trip delay experienced by the OLTP server is independent of the volume of low priority traffic or the ISL congestion. The backup traffic uses the available ISL bandwidth when it is not used by the OLTP traffic.



Tip

To achieve this traffic differentiation, be sure to enable FCC (see the [“Enabling FCC” section on page 25-3](#)).

# Configuring Data Traffic

To configure QoS, follow these steps.

- Step 1

Enable the QoS feature (see the [“Enabling QoS for Data Traffic” section on page 25-6](#)).
- Step 2

Create and define class maps (see the [“Creating Class Maps” section on page 25-7](#)).
- Step 3

Define service policies (see [“Defining Service Policies” section on page 25-8](#)).
- Step 4

Apply the configuration (see [“Applying a Service Policy” section on page 25-8](#)).

# Enabling QoS for Data Traffic

By default, the QoS data traffic feature is disabled for data traffic. To configure QoS for data traffic, you must first enable the data traffic feature in the switch.

To enable the QoS data traffic feature, follow these steps:

|        | Command                                    | Purpose                                                                                                                |
|--------|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <code>config t</code>              | Enters configuration mode.                                                                                             |
| Step 2 | switch(config)# <code>qos enable</code>    | Enables QoS. You can now configure data traffic parameters.                                                            |
|        | switch(config)# <code>no qos enable</code> | Removes the currently-applied QoS configuration and disables QoS. You can no longer configure data traffic parameters. |



Tip

QoS is supported in interoperability mode—it's effectiveness depends on the location of MDS switches in the fabric relative to the location of the source or destination of the prioritized devices.



## Creating Class Maps

Use the **class-map** option to create and define a traffic class with match criteria to identify traffic belonging to that class. Define each match criterion with one match statement from the class map configuration (`switch(config-cmap)`) mode. The class map name is restricted to 63 alphanumeric characters and defaults to the **match-all** option. Flow-based traffic uses one of the following values:

- **WWN**—Use the **source-wwn** option to specify the source WWN or the **destination-wwn** option to specify the destination WWN.
- **Fibre Channel ID (FC ID)** —Use the **source-address** option to specify the source ID (SID) or the **destination-address** option to specify the destination ID (DID). The possible values for mask are FFFFFFFF (the entire FCID is used—this is the default), FFFF00 (only domain and area FCID is used), or FF0000 (only domain FCID is used).



**Note** A **source-address** or **destination-address** of 0x000000 is not allowed.

- **Source interface**—Use the **input-interface** option to specify the ingress interface.



**Note**

The order of entries to be matched within a class map is not significant.

To create a class map, follow these steps:

|        | Command                                                                                                           | Purpose                                                                                                                                                                             |
|--------|-------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch(config)# qos class-map MyClass</code><br><code>switch(config-cmap)#</code>                           | Creates a class map called MyClass and places you in the class-map submode to match all criteria specified for this class.                                                          |
|        | <code>switch(config)# qos class-map MyClass</code><br><code>match-all</code><br><code>switch(config-cmap)#</code> | Specifies a logical AND operator for all matching statements in this class. If a frame matches all (default) configured criteria, it qualifies for this class. This is the default. |
|        | <code>switch(config)# qos class-map MyClass</code><br><code>match-any</code><br><code>switch(config-cmap)#</code> | Specifies a logical OR operator for all matching statements in this class. If a frame matches any one configured criteria, it qualifies for this class.                             |
| Step 2 | <code>switch(config-cmap)# match</code><br><code>destination-address 0x12ee00</code>                              | Specifies a destination address match for frames with the specified destination FC ID.                                                                                              |
|        | <code>switch(config-cmap)# match source-address</code><br><code>0x6dl090 mask 0xFFFFFFFF</code>                   | Specifies a source address and mask match for frames with the specified source FC ID.                                                                                               |
| Step 3 | <code>switch(config-cmap)# match destination-wwn</code><br><code>20:01:00:05:30:00:28:df</code>                   | Specifies a destination WWN to match frames.                                                                                                                                        |
|        | <code>switch(config-cmap)# match source-wwn</code><br><code>23:15:00:05:30:00:2a:1f</code>                        | Specifies a source WWN to match frames.                                                                                                                                             |
| Step 4 | <code>switch(config-cmap)# match input-interface fc</code><br><code>2/1</code>                                    | Specifies a source interface to match frames.                                                                                                                                       |
| Step 5 | <code>switch(config-cmap)# no match input-interface</code><br><code>fc 3/5</code>                                 | Removes a match based on the specified source interface.                                                                                                                            |

## Defining Service Policies

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. You can specify multiple class maps within a policy map, and map a class map to a **high**, **medium**, or **low** service level. The default priority is **low**.

As an alternative, you can map a class map to a Differentiated Services Code Point (DSCP). The DSCP is an indicator of the service level for a specified frame. The DSCP value ranges from 0 to 63, and the default is 0. A DSCP value of 46 is disallowed.

The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame.



### Note

Refer to <http://www.cisco.com/warp/public/105/dscpvalues.html#dscpandassuredforwardingclasses> for further information on implementing QoS DSCP values.

Use the **policy-map** option to specify the class of service. The policy map name is restricted to 63 alphanumeric characters.



### Note

Class-maps are processed in the order in which they are configured in each policy-map.

To specify a service policy, follow these steps:

|        | Command                                                                                   | Purpose                                                                                           |
|--------|-------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch(config)# qos policy-map MyPolicy</code><br><code>switch(config-pmap)#</code> | Creates a policy map called MyPolicy and places you in the policy-map submode.                    |
|        | <code>switch(config)# no qos policy-map OldPolicy</code><br><code>switch(config)#</code>  | Deletes the policy map called OldPolicy and places you in the policy-map submode.                 |
| Step 2 | <code>switch(config-pmap)# class MyClass</code><br><code>switch(config-pmap-c)#</code>    | Specifies the name of a predefined class and places you at the policy-map submode for that class. |
|        | <code>switch(config-pmap)# no class OldClass</code>                                       | Removes the class map called OldClass from the policy map.                                        |
| Step 3 | <code>switch(config-pmap-c)# priority high</code>                                         | Specifies the priority to be given for each frame matching this class.                            |
|        | <code>switch(config-pmap-c)# no priority high</code>                                      | Deletes a previously-assigned priority and reverts to the default value of <b>low</b> .           |
| Step 4 | <code>switch(config-pmap-c)# dscp 2</code>                                                | Specifies the DSCP value to mark each frame matching this class.                                  |
|        | <code>switch(config-pmap-c)# no dscp 60</code>                                            | Deletes a previously-assigned DSCP value and reverts to the factory default of 0.                 |

## Applying a Service Policy

When you have configured a QoS data traffic policy, you must enforce the data traffic configuration by applying that policy to the required VSAN(s). If you do not apply the policy to a VSAN, the data traffic configuration will not be enforced. You can only apply one policy map to a VSAN.

To apply a service policy, follow these steps:

|        | Command                                                             | Purpose                                                 |
|--------|---------------------------------------------------------------------|---------------------------------------------------------|
| Step 1 | <code>switch(config)# qos service policy MyPolicy vsan 3</code>     | Applies a configured policy to VSAN 3.                  |
|        | <code>switch(config)# no qos service policy OldPolicy vsan 7</code> | Deletes a configured policy that was applied to VSAN 7. |



#### Note

You can apply the same policy to a range of VSANs.

## Scheduling Traffic

The SAN-OS software supports four scheduling queues:

- Strict-priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling queues:
  - Use **dwrr-q high** option to schedule high priority traffic.
  - Use **dwrr-q medium** option to schedule medium priority traffic.
  - Use **dwrr-q low** option to schedule low priority traffic.

Use the **qos dwrr-q** command to associate a weight with a DWRR queue.

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for high queue, 30 for the medium queue, and 20 for the low queue. Decreasing order of queue weights is mandated to ensure the higher priority queues have a higher service level, though the ratio of the configured weights can vary (for example, one can configure 70:30:5 or 60:50:10 but not 50:70:10).

To associate a weight with a DWRR queue, follow these steps:

|        | Command                                                | Purpose                                                                 |
|--------|--------------------------------------------------------|-------------------------------------------------------------------------|
| Step 1 | <code>switch(config)# qos dwrr-q high weight 10</code> | Associates a relative weight (10) to a specified queue (default queue). |
|        | <code>switch(config)# no dwrr-q low weight 51</code>   | Restores the default weight of 20.                                      |

## Displaying Data Traffic Information

The **show qos** commands display the current QoS settings for data traffic (see Examples 25-3 to 25-3).

### Example 25-3 Displays the Contents of all Class Maps

```
switch# show qos class-map
qos class-map MyClass match-any
 match destination-wwn 20:01:00:05:30:00:28:df
 match source-wwn 23:15:00:05:30:00:2a:1f
 match input-interface fc2/1
qos class-map Class2 match-all
 match input-interface fc2/14
```

```
qos class-map Class3 match-all
 match source-wnn 20:01:00:05:30:00:2a:1f
```

The **show qos** command

**Example 25-4** *Displays the Contents of a Specified Class Map*

```
switch# show qos class-map name MyClass
qos class-map MyClass match-any
 match destination-wnn 20:01:00:05:30:00:28:df
 match source-wnn 23:15:00:05:30:00:2a:1f
 match input-interface fc2/1
```

**Example 25-5** *Displays All Configured Policy Maps*

```
switch# show qos policy-map
qos policy-map MyPolicy
 class MyClass
 priority medium

qos policy-map Policy1
 class Class2
 priority low
```

**Example 25-6** *Displays a Specified Policy Map*

```
switch# show qos policy-map name MyPolicy
qos policy-map MyPolicy
 class MyClass
 priority medium
```

**Example 25-7** *Displays Scheduled DWRR Configurations*

```
switch# show qos dwrr
qos dwrr-q high weight 50
qos dwrr-q medium weight 30
qos dwrr-q low weight 20
```

**Example 25-8** *Displays All Applied Policy Maps*

```
switch# show qos service policy
qos service policy MyPolicy vsan 1
qos service policy Policy1 vsan 4
```

**Example 25-9** *Displays the Policy Map Associated with a Specified VSAN*

```
switch# show qos service policy vsan 1
qos policy-map pmap1
 class cmap1
 priority medium
 class cmap2
 priority high
```

**Example 25-10** *Displays the Class Map Associated with a Specified Interface*

```
switch# show qos service policy interface fc3/10
qos policy-map pmap1
 class cmap3
 priority high
 class cmap4
```

```
priority low
```

***Example 25-11 Displays QoS Statistics***

```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 301431
Number of highest-priority FC frames transmitted = 137679
Current priority of FC control frames = 7 (0 = lowest; 7 = highest)
```

# Ingress Port Rate Limiting

A port rate limiting feature is available in SAN-OS 1.3(x). This feature helps control the bandwidth for individual FC ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a FC port. The feature controls traffic flow by limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports.



## Note

Port rate limiting can only be configured in switches in the Cisco MDS 9100 Series.

This command can only be configured if the following conditions hold true:

- The QoS feature is enabled using the **qos enable** command.
- The command is issued in a Cisco MDS 9100 series switch.

The rate limit ranges from 1 to 100% and the default is 100%.

To configure the port rate limiting value, follow these steps.

|        | Command                                                 | Purpose                                                              |
|--------|---------------------------------------------------------|----------------------------------------------------------------------|
| Step 1 | switch # <b>config t</b><br>switch(config)#             | Enters the configuration mode.                                       |
| Step 2 | switch(config)# <b>interface fc 1/1</b>                 | Selects the interface to specify the ingress port rate limit.        |
| Step 3 | switch(config-if)# <b>switchport ingress-rate 50</b>    | Configures a 50% port rate limit for the selected interface.         |
|        | switch(config-if)# <b>no switchport ingress-rate 50</b> | Reverts a previously-configured rate to the factory default of 100%. |

## Default Settings

Table 25-1 lists the default settings for FCC, QoS, and rate limiting features:

**Table 25-1 Default FCC, QoS, and Rate Limiting Settings**

| Parameters          | Default   |
|---------------------|-----------|
| FCC protocol        | Disabled. |
| QoS control traffic | Enabled.  |
| QoS data traffic    | Disabled. |
| Rate limit          | 100%      |



## Configuring System Message Logging

---

This chapter describes how to configure system message logging on the Cisco MDS 9000 Family switches. It includes the following sections:

- [About System Message Logging, page 26-2](#)
- [Configuring System Message Logging, page 26-4](#)
- [Displaying System Message Logging Information, page 26-8](#)
- [Default Settings, page 26-12](#)

# About System Message Logging

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting
- Allows you to select the types of captured logging information.
- Allows you to select the destination of the captured logging information.

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility (see [Table 26-1](#)) and the severity level (see [Table 26-2](#)). Messages are time-stamped to enhance real-time debugging and management.

You can access logged system messages using the CLI or by saving them to a properly configured syslog server. The switch software saves syslog messages in a file that can be configured to save up to 4 MB. You can monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a syslog server.



## Note

When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a syslog server for a few seconds.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM. You can view this log at any time using the **show logging nvram** command.

[Table 26-1](#) describes the facilities supported by the system message logs.

**Table 26-1 Internal Logging Facilities**

| Facility Keyword | Description                    | Standard or Cisco MDS Specific |
|------------------|--------------------------------|--------------------------------|
| <b>acl</b>       | ACL manager                    | Cisco MDS 9000 Family specific |
| <b>all</b>       | All facilities                 | Cisco MDS 9000 Family specific |
| <b>auth</b>      | Authorization system           | Standard                       |
| <b>authpriv</b>  | Authorization (private) system | Standard                       |
| <b>bootvar</b>   | Bootvar                        | Cisco MDS 9000 Family specific |
| <b>callhome</b>  | Call Home                      | Cisco MDS 9000 Family specific |
| <b>cron</b>      | Cron or at facility            | Standard                       |
| <b>daemon</b>    | System daemons                 | Standard                       |
| <b>fcc</b>       | FCC                            | Cisco MDS 9000 Family specific |
| <b>fcdomain</b>  | fcdomain                       | Cisco MDS 9000 Family specific |
| <b>fcns</b>      | Name server                    | Cisco MDS 9000 Family specific |
| <b>fcs</b>       | FCS                            | Cisco MDS 9000 Family specific |
| <b>flogi</b>     | FLOGI                          | Cisco MDS 9000 Family specific |
| <b>fspf</b>      | FSPF                           | Cisco MDS 9000 Family specific |
| <b>ftp</b>       | File Transfer Protocol         | Standard                       |



**Table 26-1 Internal Logging Facilities (continued)**

| Facility Keyword        | Description                      | Standard or Cisco MDS Specific |
|-------------------------|----------------------------------|--------------------------------|
| <b>ipconf</b>           | IP configuration                 | Cisco MDS 9000 Family specific |
| <b>ipfc</b>             | IPFC                             | Cisco MDS 9000 Family specific |
| <b>kernel</b>           | Kernel                           | Standard                       |
| <b>local0 to local7</b> | Locally defined messages         | Standard                       |
| <b>lpr</b>              | Line printer system              | Standard                       |
| <b>mail</b>             | Mail system                      | Standard                       |
| <b>mcast</b>            | Multicast                        | Cisco MDS 9000 Family specific |
| <b>module</b>           | Switching module                 | Cisco MDS 9000 Family specific |
| <b>news</b>             | USENET news                      | Standard                       |
| <b>ntp</b>              | NTP                              | Cisco MDS 9000 Family specific |
| <b>platform</b>         | Platform manager                 | Cisco MDS 9000 Family specific |
| <b>port</b>             | Port                             | Cisco MDS 9000 Family specific |
| <b>port-channel</b>     | PortChannel                      | Cisco MDS 9000 Family specific |
| <b>qos</b>              | QoS                              | Cisco MDS 9000 Family specific |
| <b>rdl</b>              | RDL                              | Cisco MDS 9000 Family specific |
| <b>rib</b>              | RIB                              | Cisco MDS 9000 Family specific |
| <b>rscn</b>             | RSCN                             | Cisco MDS 9000 Family specific |
| <b>securityd</b>        | Security                         | Cisco MDS 9000 Family specific |
| <b>syslog</b>           | Internal syslog messages         | Standard                       |
| <b>sysmgr</b>           | System manager                   | Cisco MDS 9000 Family specific |
| <b>tlport</b>           | TL port                          | Cisco MDS 9000 Family specific |
| <b>user</b>             | User process                     | Standard                       |
| <b>uucp</b>             | Unix-to-Unix copy system         | Standard                       |
| <b>vhbad</b>            | Virtual host base adapter daemon | Cisco MDS 9000 Family specific |
| <b>vni</b>              | Virtual network interface        | Cisco MDS 9000 Family specific |
| <b>vrrp_cfg</b>         | VRRP configuration               | Cisco MDS 9000 Family specific |
| <b>vrrp_eng</b>         | VRRP engine                      | Cisco MDS 9000 Family specific |
| <b>vsan</b>             | VSAN syslog                      | Cisco MDS 9000 Family specific |
| <b>vshd</b>             | vshd                             | Cisco MDS 9000 Family specific |
| <b>wwn</b>              | WWN manager                      | Cisco MDS 9000 Family specific |
| <b>xbar</b>             | Xbar syslog                      | Cisco MDS 9000 Family specific |
| <b>zone</b>             | Zone server                      | Cisco MDS 9000 Family specific |

Table 26-2 describes the severity levels supported by the system message logs.

**Table 26-2 Error Message Severity Levels**

| Level Keyword | Level | Description                      | Syslog Definition |
|---------------|-------|----------------------------------|-------------------|
| emergencies   | 0     | System unusable                  | LOG_EMERG         |
| alerts        | 1     | Immediate action needed          | LOG_ALERT         |
| critical      | 2     | Critical conditions              | LOG_CRIT          |
| errors        | 3     | Error conditions                 | LOG_ERR           |
| warnings      | 4     | Warning conditions               | LOG_WARNING       |
| notifications | 5     | Normal but significant condition | LOG_NOTICE        |
| informational | 6     | Informational messages only      | LOG_INFO          |
| debugging     | 7     | Debugging messages               | LOG_DEBUG         |



**Note**

Refer to the *Cisco MDS 9000 Family System Messages Guide* for details on the error log message format.

## Configuring System Message Logging

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

### Enabling Message Logging

You can disable logging to the console or enable logging to a given Telnet or SSH session.

- When you disable or enable logging to a console session, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved.
- When you enable or disable logging to a Telnet or SSH session, that state is applied only to that session. If you exit and log in again to a new session, the state is not preserved.

To enable or disable the logging state for a Telnet, or SSH session, follow these steps:

|        | Command                            | Purpose                                                                                                       |
|--------|------------------------------------|---------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>terminal monitor</b>    | Enables logging for a Telnet, or SSH session.<br><b>Note</b> A console session is enabled by default.         |
| Step 2 | switch# <b>terminal no monitor</b> | Disables logging for a Telnet, or SSH session.<br><b>Note</b> A Telnet or SSH session is disabled by default. |

## Configuring Console Severity Level

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is 2 (critical).

To configure the severity level for a logging facility, follow these steps:

|        | Command                                    | Purpose                                                                                                                                                                   |
|--------|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)# | Enters configuration mode.                                                                                                                                                |
| Step 2 | switch(config)# <b>logging console 3</b>   | Configures console logging at level 3 (error). Logging messages with a severity level of 3 or above will be displayed on the console.                                     |
|        | switch(config)# <b>logging console</b>     | Reverts console logging to the factory set default severity level of 2 (critical). Logging messages with a severity level of 2 or above will be displayed on the console. |



### Tip

The current critical (default) logging level is maintained, if the console baud speed is 9600 baud (default). All attempts to change the console logging level generates an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud (see the [“Configuring Console Settings”](#) section on page 4-31).

## Configuring Module Logging

By default, logging is enabled at Level 7 for all modules. You can enable or disable logging for each module at a specified level.

To configure the severity level for a logging facility, follow these steps:

|        | Command                                    | Purpose                                                                                                                                                                       |
|--------|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)# | Enters configuration mode.                                                                                                                                                    |
| Step 2 | switch(config)# <b>logging module 1</b>    | Configures module logging at Level 1 (alerts).                                                                                                                                |
|        | switch(config)# <b>logging module</b>      | Configures module logging for all modules in the switch.                                                                                                                      |
|        | switch(config)# <b>no logging console</b>  | Reverts console logging to the factory set default severity level of 5 (notification). Logging messages with a severity level of 5 or above will be displayed on the console. |

## Configuring Facility Severity Level

To configure the severity level for a logging facility, follow these steps:

|        | Command                                       | Purpose                                                                                                                                                             |
|--------|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | switch# <b>config t</b><br>switch(config)#    | Enters configuration mode.                                                                                                                                          |
| Step 2 | switch(config)# <b>logging level kernel 4</b> | Configures Telnet or SSH logging for the kernel facility at level 4 (warning). As a result, logging messages with a severity level of 4 or above will be displayed. |

## Configuring Log Files

Logging messages may be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is `messages`. You can rename this file using the **logging logfile** command. The file name can have up to 200 characters and the file size ranges from 4096 bytes to 4194304 bytes.

To send log messages to file, follow these steps:

|        | Command                                                                                | Purpose                                                                                                                                                                                                                                                    |
|--------|----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <code>switch# config t</code><br><code>switch(config)#</code>                          | Enters configuration mode.                                                                                                                                                                                                                                 |
| Step 2 | <code>switch(config)# logging logfile</code><br><code>ManagerLog 3 size 3000000</code> | Configures logging information for errors or events above severity level 3 to be logged in a file named <code>ManagerLog</code> . By configuring a size, you are restricting the file size to 3000000 bytes. The maximum upper limit is 4194304 (default). |

The configured log file is saved in the `/var/log/external` directory. The location of the log file cannot be changed. You can use the **show logging** and **clear debug-logfile** commands to view and clear this file. It is not accessible using the **dir** command.

You can display the log file using the **show logging logfile** command and copy the logfile to a different location using the **copy log:messages** command using additional copy syntax (see the [“Copying Files” section on page 4-27](#)).

## Configuring Syslog Servers

To send log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. Log in as root, and perform these steps:

- Step 1** Add the following line to the file `/etc/syslog.conf`
- ```
local1.debug                /var/log/myfile.log
```



Note Be sure to add five tab characters between **local1.debug** and **/var/log/myfile.log**. Refer to entries in the `/etc/syslog.conf` file for further examples.

The switch sends messages according to the specified facility types and severity levels. The **local1** keyword specifies the UNIX logging facility used. The messages from the switch are generated by user processes. The **debug** keyword specifies the severity level of the condition being logged. You can set UNIX systems to receive all messages from the switch.

- Step 2** Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/myfile.log
$ chmod 666 /var/log/myfile.log
```

- Step 3** Make sure the syslog daemon reads the new changes by entering this command:

```
$ kill -HUP ~cat /etc/syslog.pid~
```

To configure syslog servers, follow these steps:

	Command	Purpose
Step 1	switch# config t switch#	Enters configuration mode.
Step 2	switch(config)# logging server 172.22.00.00 switch(config)#	Configures the switch to forward log messages according to the specified facility types and severity levels to remote multiple servers specified by its hostname or IP address (172.22.00.00). Note You can configure a maximum of three syslog servers.
	switch(config)# logging server 172.22.00.00 facility local1 switch(config)#	Configures the switch to forward log messages according to the specified facility (local1) for the server IP address (172.22.00.00). The default outgoing facility is local7.
	switch(config)# no logging server 172.11.00.00 switch(config)#	Removes the specified server (172.11.00.00) and reverts to factory default. Note You can configure a maximum of three syslog servers.

Outgoing Syslog Server Logging Facilities

All syslog messages have a logging facility and a level. The logging facility can be thought of as where and the level can be thought of as what.

The single syslog daemon (syslogd) sends the information based on the configured **facility** option. If no facility is specified, local7 is the default outgoing facility.

The internal facilities are listed in [Table 26-1](#) and the outgoing logging facilities are listed in [Table 26-3](#).

Table 26-3 Outgoing Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
cron	Cron or at facility	Standard
daemon	System daemons	Standard
ftp	File Transfer Protocol	Standard
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard (local7 is the default)
lpr	Line printer system	Standard
mail	Mail system	Standard
news	USENET news	Standard
syslog	Internal syslog messages	Standard
user	User process	Standard
uucp	Unix-to-Unix copy system	Standard

Displaying System Message Logging Information

Use the **show logging** command to display the current system message logging configuration. See Examples 26-1 to 26-10.

Example 26-1 Displays Current System Message Logging

```
switch# show logging
Logging console:          enabled (Severity: critical)
Logging monitor:          enabled (Severity: debugging)
Logging linecard:         enabled (Severity: debugging)
Logging server:           enabled
{172.20.102.34}
    server severity:       debugging
    server facility:       local7
{10.77.202.88}
    server severity:       debugging
    server facility:       local7
{10.77.202.149}
    server severity:       debugging
    server facility:       local7
Logging logfile:          enabled
Name - messages: Severity - debugging Size - 4194304
```

Facility	Default Severity	Current Session Severity
-----	-----	-----
kern	6	6
user	3	3
mail	3	3
daemon	7	7
auth	0	7
syslog	3	3
lpr	3	3
news	3	3
uucp	3	3
cron	3	3
authpriv	3	7
ftp	3	3
local0	3	3
local1	3	3
local2	3	3
local3	3	3
local4	3	3
local5	3	3
local6	3	3
local7	3	3
vsan	2	2
fspf	3	3
fcdomain	2	2
module	5	5
sysmgr	3	3
zone	2	2
vni	2	2
ipconf	2	2
ipfc	2	2
xbar	3	3
fcns	2	2
fcs	2	2
acl	2	2
tlport	2	2
port	5	5
flogi	2	2

```

port_channel      5          5
wnn                3          3
fcc                2          2
qos                3          3
vrrp_cfg          2          2
ntp                2          2
platform          5          5
vrrp_eng           2          2
callhome           2          2
mcast              2          2
rdl                2          2
rscn               2          2
bootvar           5          2
securityd          2          2
vhbad              2          2
rib                2          2
vshd               5          5

0(emergencies)     1(alerts)     2(critical)
3(errors)          4(warnings)  5(notifications)
6(information)    7(debugging)

```

```

Feb 14 09:50:57 excal-113 %TTYD-6-TTYD_MISC: TTYD TTYD started
Feb 14 09:50:58 excal-113 %DAEMON-6-SYSTEM_MSG: precision = 8 usec
...

```

Use the **show logging nvram** command to view the log messages saved in NVRAM. Only log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

Example 26-2 Displays NVRM Log Contents

```

switch# show logging nvram
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2209, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2199, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
...

```

Example 26-3 Displays the Log File

```

switch# show logging logfile
Jul 16 21:06:50 %DAEMON-3-SYSTEM_MSG: Un-parsable frequency in /mnt/pss/ntp.drift
Jul 16 21:06:56 %DAEMON-3-SYSTEM_MSG: snmpd:snmp_open_debug_cfg: no snmp_saved_dbg_uri ;
Jul 16 21:06:58 172.22.91.204 %PORT-5-IF_UP: Interface mgmt0 is up
Jul 16 21:06:58 172.22.91.204 %MODULE-5-ACTIVE_SUP_OK: Supervisor 5 is active
...

```

Example 26-4 Displays Console Logging Status

```

switch# show logging console
Logging console:                  enabled (Severity: notifications)

```

Example 26-5 Displays Logging Facility

```

switch# show logging level
Facility      Default Severity  Current Session Severity
-----
kern          6                  6

```

■ Displaying System Message Logging Information

user	3	3
mail	3	3
daemon	7	7
auth	0	7
syslog	3	3
lpr	3	3
news	3	3
uucp	3	3
cron	3	3
authpriv	3	7
ftp	3	3
local0	3	3
local1	3	3
local2	3	3
local3	3	3
local4	3	3
local5	3	3
local6	3	3
local7	3	3
vsan	2	2
fspf	3	3
fcdomain	2	2
module	5	5
sysmgr	3	3
zone	2	2
vni	2	2
ipconf	2	2
ipfc	2	2
xbar	3	3
fcns	2	2
fcs	2	2
acl	2	2
tlport	2	2
port	5	5
flogi	2	2
port_channel	5	5
wwn	3	3
fcc	2	2
qos	3	3
vrrp_cfg	2	2
ntp	2	2
platform	5	5
vrrp_eng	2	2
callhome	2	2
mcast	2	2
rdl	2	2
rscn	2	2
bootvar	5	2
securityd	2	2
vhbad	2	2
rib	2	2
vshd	5	5
0 (emergencies)	1 (alerts)	2 (critical)
3 (errors)	4 (warnings)	5 (notifications)
6 (information)	7 (debugging)	

Example 26-6 Displays Logging Information

```

switch# show logging info
Logging console:          enabled (Severity: critical)
Logging monitor:          enabled (Severity: debugging)
Logging linecard:         enabled (Severity: debugging)

```



```

Logging server:                enabled
{172.20.102.34}
    server severity:           debugging
    server facility:           local7
{10.77.202.88}
    server severity:           debugging
    server facility:           local7
{10.77.202.149}
    server severity:           debugging
    server facility:           local7
Logging logfile:               enabled
    Name - messages: Severity - debugging Size - 4194304

```

Facility -----	Default Severity -----	Current Session Severity -----
kern	6	6
user	3	3
mail	3	3
daemon	7	7
auth	0	7
syslog	3	3
lpr	3	3
news	3	3
uucp	3	3
cron	3	3
authpriv	3	7
ftp	3	3
local0	3	3
local1	3	3
local2	3	3
local3	3	3
local4	3	3
local5	3	3
local6	3	3
local7	3	3
vsan	2	2
fspf	3	3
fcdomain	2	2
module	5	5
sysmgr	3	3
zone	2	2
vni	2	2
ipconf	2	2
ipfc	2	2
xbar	3	3
fcns	2	2
fcs	2	2
acl	2	2
tlport	2	2
port	5	5
flogi	2	2
port_channel	5	5
wwn	3	3
fcc	2	2
qos	3	3
vrrp_cfg	2	2
ntp	2	2
platform	5	5
vrrp_eng	2	2
callhome	2	2
mcast	2	2
rdl	2	2
rscn	2	2
bootvar	5	2

```

securityd          2          2
vhbad              2          2
rib                2          2
vshd               5          5

0(emergencies)     1(alerts)    2(critical)
3(errors)          4(warnings)  5(notifications)
6(information)     7(debugging)

```

Example 26-7 Displays Last Few Lines of a Log File

```

switch# show logging last 2
Nov 8 16:48:04 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/1
(171.71.58.56)
Nov 8 17:44:09 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
(171.71.58.72)

```

**Note**

Use the **show logging filename** command to display the entire log file.

Example 26-8 Displays Switching Module Logging Status

```

switch# show logging module
Logging linecard:          enabled (Severity: debugging)

```

Example 26-9 Displays Monitor Logging Status

```

switch# show logging monitor
Logging monitor:          enabled (Severity: information)

```

**Note**

Use the **show logging nvram** command to view the log messages in NVRAM.

Example 26-10 Displays Server Information

```

switch# show logging server
Logging server:          enabled
{172.22.95.167}
    server severity:      debugging
    server facility:      local7
{172.22.92.58}
    server severity:      debugging
    server facility:      local7

```

Default Settings

Table 26-4 lists the default settings for system message logging.

Table 26-4 Default System Message Log Setting

Parameters	Default
System message logging to the console	Enabled for messages at the critical severity level.
System message logging to Telnet sessions	Disabled.
Logging file size	4194304.

Table 26-4 Default System Message Log Setting (continued)

Parameters	Default
Log file name	message (can be changed to any name with up to 200 characters).
Logging server	Disabled.
Syslog server IP address	Non configured.
No. of servers	3 servers.
Server facility	Local 7.



Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [About SCSI LUN Discovery, page 27-1](#)
- [Starting SCSI LUN Discovery, page 27-2](#)
- [Initiating Customized Discovery, page 27-2](#)
- [Displaying SCSI LUN Information, page 27-3](#)

About SCSI LUN Discovery

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server.

The name server requires LUN information for the following reasons:

- To display LUN storage device information so an NMS can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco MDS 9000 Family.

Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

To begin SCSI LUN discovery, follow this step:

	Command	Purpose
Step 1	switch# discover scsi-target local os all discovery started	Discovers local SCSI targets for all Operating Systems (OS). The operating system options are aix , all , hpux , linux , solaris , or windows
	switch# discover scsi-target remote os aix discovery started	Discovers remote SCSI targets assigned to the AIX OS.
	switch# discover scsi-target vsan 1 fcid 0x9c03d6 discover scsi-target vsan 1 fcid 0x9c03d6 VSAN: 1 FCID: 0x9c03d6 PWWN: 00:00:00:00:00:00:00:00 PRLI RSP: 0x01 SPARM: 0x0012 SCSI TYPE: 0 NLUNS: 1 Vendor: Company 4 Model: ST318203FC Rev: 0004 Other: 00:00:02:32:8b:00:50:0a	Discovers SCSI targets for the specified VSAN (1) and FC ID (0x9c03d6).
	switch# discover scsi-target custom-list os linux discovery started	Discovers SCSI targets from the customized list assigned to the Linux OS.

Only Nx ports present in the name server database and which have registered a FC4 Type = SCSI_FCP are discovered.

Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. Use the **custom-list** option to initiate this discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

To initiate a customized discovery, follow this step:

	Command	Purpose
Step 1	switch# discover custom-list add vsan 1 domain 0X123456	Adds the specified entry to the custom list.
	switch# discover custom-list delete vsan 1 domain 0X123456	Deletes the specified domain ID from the custom list.

Displaying SCSI LUN Information

Use the **show scsi-target** and **show fcns database** commands to display the results of the discovery. See Examples 27-1 to 27-7.

Example 27-1 Displays the Discovered Targets

```
switch# show scsi-target status
discovery completed
```



Note

The discovery can take several minutes to complete, especially if the fabric is large fabric or if several devices are slow to respond.

Example 27-2 Displays the FCNS Database

```
switch# show fcns database
172.22.91.115# show fcns database
```

VSAN 1:

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0xeb0000	N	21:01:00:e0:8b:2a:f6:54	(Qlogic)	scsi-fcp:init
0xeb0201	NL	10:00:00:00:c9:32:8d:76	(Emulex)	scsi-fcp:init

Total number of entries = 2

VSAN 7:

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0xed0001	NL	21:00:00:04:cf:fb:42:f8	(Seagate)	scsi-fcp:target

Total number of entries = 1

VSAN 2002:

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0xcafe00	N	20:03:00:05:30:00:2a:20	(Cisco)	FICON:CUP

Total number of entries = 1

Example 27-3 Displays the Discovered Target Disks

```
switch# show scsi-target disk
```

VSAN	FCID	PWWN	VENDOR	MODEL	REV
1	0x9c03d6	21:00:00:20:37:46:78:97	Company 4	ST318203FC	0004
1	0x9c03d9	21:00:00:20:37:5b:cf:b9	Company 4	ST318203FC	0004
1	0x9c03da	21:00:00:20:37:18:6f:90	Company 4	ST318203FC	0004
1	0x9c03dc	21:00:00:20:37:5a:5b:27	Company 4	ST318203FC	0004
1	0x9c03e0	21:00:00:20:37:36:0b:4d	Company 4	ST318203FC	0004
1	0x9c03e1	21:00:00:20:37:39:90:6a	Company 4	ST318203 CLAR18	3844
1	0x9c03e2	21:00:00:20:37:18:d2:45	Company 4	ST318203 CLAR18	3844
1	0x9c03e4	21:00:00:20:37:6b:d7:18	Company 4	ST318203 CLAR18	3844
1	0x9c03e8	21:00:00:20:37:38:a7:c1	Company 4	ST318203FC	0004
1	0x9c03ef	21:00:00:20:37:18:17:d2	Company 4	ST318203FC	0004

Example 27-4 Displays the Discovered LUNs for All OSs

```
switch# show scsi-target lun os all

ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
-----
OS  LUN      Capacity Status  Serial Number  Device-Id
      (MB)
-----
WIN 0x0      36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
AIX 0x0      36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
SOL 0x0      36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
LIN 0x0      36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
HP  0x0      36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

Example 27-5 Displays the Discovered LUNs for the Solaris OS

```
switch# show scsi-target lun os solaris

ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
-----
OS  LUN      Capacity Status  Serial Number  Device-Id
      (MB)
-----
SOL 0x0      36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

The following command displays the port WWN that is assigned to each OS (Windows, AIX, Solaris, Linux, or HPUX)

Example 27-6 Displays the pWWNs for each OS

```
switch# show scsi-target pwwn
-----
OS      PWWN
-----
WIN     24:91:00:05:30:00:2a:1e
AIX     24:92:00:05:30:00:2a:1e
SOL     24:93:00:05:30:00:2a:1e
LIN     24:94:00:05:30:00:2a:1e
HP      24:95:00:05:30:00:2a:1e
```

Example 27-7 Displays Customized Discovered Targets

```
switch# show scsi-target custom-list
-----
VSAN    DOMAIN
-----
1       56
```

Use the **show scsi-target auto-poll** command to verify automatic discovery of scsi-targets which come online. The internal uuid number indicates that a CSM or an IPS module is in the chassis.

Example 27-8 Displays Customized Discovered Targets

```
switch# show scsi-target auto-poll
auto-polling is enabled, poll_start:0 poll_count:1 poll_type:0
USERS OF AUTO POLLING
```

uuid:54



Monitoring Network Traffic Using SPAN

This chapter describes the switched port analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

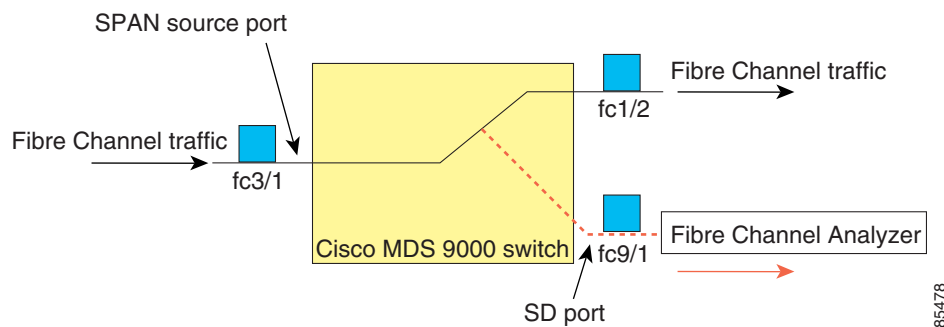
- [About SPAN, page 28-2](#)
- [SPAN Sources, page 28-2](#)
- [SPAN Sessions, page 28-5](#)
- [Specifying Filters, page 28-5](#)
- [SD Port Characteristics, page 28-6](#)
- [Configuring SPAN, page 28-7](#)
- [Monitoring Traffic Using Fibre Channel Analyzers, page 28-10](#)
- [Displaying SPAN Information, page 28-13](#)
- [Default SPAN Settings, page 28-14](#)
- [Remote SPAN, page 28-15](#)

About SPAN

The switched port analyzer (SPAN) feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD-port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic (see “[Configuring a Fabric Analyzer](#)” section on page 29-7).

SD ports do not receive frames, they merely transmit a copy of the SPAN source traffic. The SPAN feature is non-intrusive and does not affect switching of network traffic for any SPAN source ports (see [Figure 28-1](#)).

Figure 28-1 SPAN Transmission

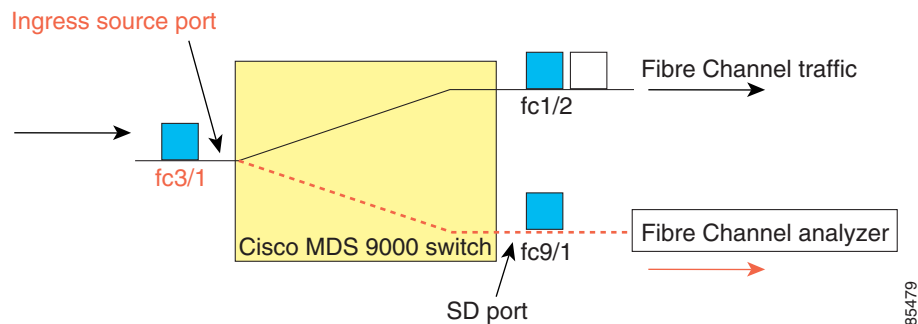


SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:

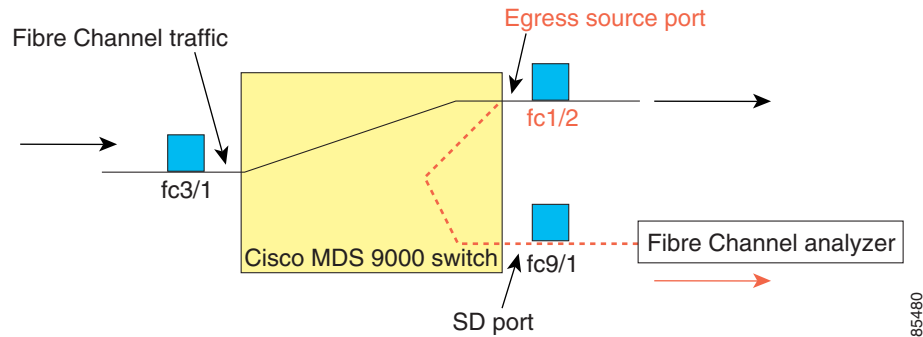
- Ingress source (rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see [Figure 28-2](#)).

Figure 28-2 SPAN Traffic from the Ingress Direction



- Egress source (tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see Figure 28-3).

Figure 28-3 SPAN Traffic from Egress Direction



IPS Source Ports

Effective SAN-OS Release 1.3(x) Switched Port Analyzer (SPAN) capabilities are also available on the IP Storage Services (IPS) module. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel port interfaces, not the physical Gigabit Ethernet ports. You can SPAN ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.



Note

You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

CSM Source Ports

Effective SAN-OS Release 1.3(x) Switched Port Analyzer (SPAN) capabilities are also available on the Caching Services Module (CSM).

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for further information.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports:
 - F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
 - The Fibre Channel traffic from the supervisor module to the switch fabric, through the sup-fc0 interface, is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
 - The Fibre Channel traffic from the switch fabric to the supervisor module, through the sup-fc0 interface, is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels

- All ports in the PortChannel are included and spanned as sources.
- You cannot specify individual ports in a PortChannel as SPAN sources. Previously-configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces.
 - iSCSI interface
 - FCIP interfaces

VSAN as a SPAN Source

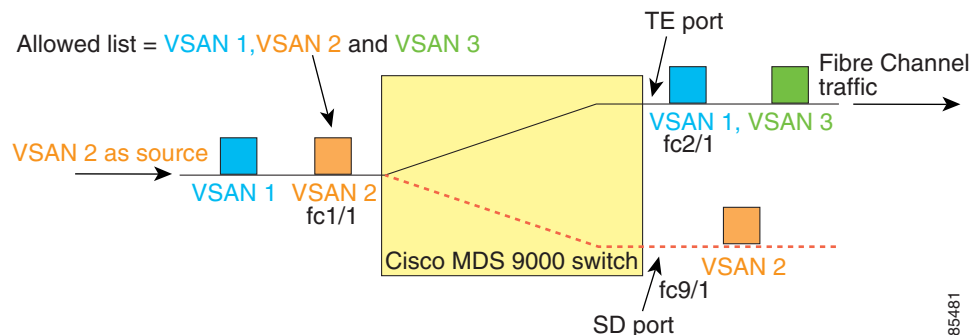
When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

Guidelines to Configure VSANs as a Source

The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- When a VSAN is specified as a source, you will not be able to perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously-configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a SPAN source, you will not be able to configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.
- Interfaces are only included as sources when the port VSAN matches the source VSAN. [Figure 28-4](#) displays a configuration using VSAN 2 as a SPAN source:
 - All ports in the switch are in VSAN 1 except fc1/1.
 - Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
 - VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 28-4 VSAN As a SPAN Source

For this configuration, the following apply:

- VSAN 2 as a SPAN source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 as the port VSAN does not match VSAN 1. See [“Configuring Trunk-Allowed VSAN List”](#) section on page 11-4 or [“VSAN Membership”](#) section on page 9-5.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate a SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic will not be directed to the SD port.

To temporarily deactivate (suspend) a SPAN session use the **suspend** command in the SPAN submode. The traffic monitoring is stopped during this time. You can reactivate the SPAN session using the **no suspend** command.



Tip

A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session (see [Figure 28-4](#)). Only VSANs present in the filter are spanned.

You can specify session VSAN filters which are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores buffer-to-buffer credits.
- Allows data traffic only in the egress (tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The port mode can not be changed if it is being used for a SPAN session.

**Note**

If you need to change a SD-port mode to another port mode, first remove the SD port from all sessions and then change the port mode using the **switchport mode** command.

- The outgoing frames can be encapsulated in extended inter-switch link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Advanced Services Modules (ASMs).

Guidelines to Configure SPAN

The following guidelines apply for a SPAN configuration:

- You can configure up to 16 SPAN sessions with multiple ingress (rx) sources.
- You can configure a maximum of three SPAN sessions with one egress (tx) port.
- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit (see the [“Configuring 32-port Switching Modules and Host-Optimized Ports”](#) section on page 10-8).
- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.

Configuring SPAN

To monitor network traffic using SD ports, follow these steps:

- Step 1** Configure the SD port.
- Step 2** Attach the SD port to a SPAN session.
- Step 3** Monitor network traffic by adding source interfaces to the session.

To configure an SD port for SPAN monitoring, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc9/1	Configures the specified interface.
Step 3	switch(config-if)# switchport mode SD	Configures the SD port-mode for interface fc2/1.
Step 4	switch(config-if)# switchport speed 1000	Configures the SD port speed to 1000 Mbps.
Step 5	switch(config-if)# no shutdown	Enables traffic flow through this interface.

To configure a SPAN session, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# span session 1 switch(config-span)#	Configures the specified SPAN session (1). If the session does not exist, it will be created.
	switch(config)# no span session 1	Deletes the specified SPAN session (1).
Step 3	switch(config-span)# destination interface fc9/1	Configures the specified destination interface (fc 9/1) in a session.
	switch(config-span)# no destination interface fc9/1	Removes the specified destination interface (fc 9/1).
Step 4	switch(config-span)# source interface fc7/1	Configures the source (fc7/1) interface in both directions.
	switch(config-span)# no source interface fc7/1	Removes the specified destination interface (fc 7/1) from this session.

	Command	Purpose
Step 5	switch(config-span)# source interface sup-fc0	Configures the source interface (sup-fc0) in the session.
	switch(config-span)# source interface fc1/5 - 6, fc2/1 -3	Configures the specified interface ranges in the session.
	switch(config-span)# source vsan 1-2	Configures source VSANs 1 and 2 in the session.
	switch(config-span)# source interface port-channel 1	Configures the source PortChannel (port-channel 1).
	switch(config-span)# source interface fcip 51	Configures the source FCIP interface in the session.
	switch(config-span)# source interface iscsi 4/1	Configures the source iSCSI interface in the session.
	switch(config-span)# source interface svc1/1 tx traffic-type initiator	Configures the source SVC interface in the Tx direction for a initiator traffic type.
	switch(config-span)# no source interface port-channel 1	Deletes the specified source interface (port-channel 1)
Step 6	switch(config-span)# suspend	Suspends the session.
	switch(config-span)# no suspend	Reactivates the session.

To configure a SPAN filter, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# span session 1 switch(config-span)#	Configures the specified session (1).
Step 3	switch(config-span)# source interface fc9/1 tx	Configures the source fc9/1 interface in the egress (tx) direction
	switch(config-span)# source filter vsan 1-2	Configures VSANs 1 and 2 as session filters.
	switch(config-span)# source interface fc7/1 rx	Configures the source fc7/1 interface in the ingress (rx) direction.

Encapsulating Frames

The **switchport encap eisl** command only applies to SD port interfaces. This command is disabled by default. If you enable the encapsulation feature, all outgoing frames will be encapsulated. If encapsulation is enabled, you will see a new line (Encapsulation is eisl) in the **show interface SD_port_interface** command output.

To encapsulate outgoing frames (optional), follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# interface fc9/32	Configures the specified interface.
Step 3	switch(config-if)# switchport mode SD	Configures the SD port-mode for interface fc2/1.

	Command	Purpose
Step 4	switch(config-if)# switchport encap eisl	Enables the encapsulation option for this SD port.
	switch(config-if)# no switchport encap eisl	Disables the encapsulation option and reverts the switch to factory default.

SPAN Conversion Behavior

Effective Release 1.1(1), SPAN features (configured in any prior release) are converted as stated below:

- If source interfaces and source VSANs are configured in a given session, then all the source VSANs are removed from that session. For example,

Before Release 1.0(4)

```
Session 1 (active)
  Destination is fc1/9
  No session filters configured
  Ingress (rx) sources are
    vsans 10-11
    fc1/3,
  Egress (tx) sources are
    fc1/3,
```

Once upgraded to Release 1.1(1):

```
Session 1 (active)
  Destination is fc1/9
  No session filters configured
  Ingress (rx) sources are
    fc1/3,
  Egress (tx) sources are
    fc1/3,
```

Session 1 had both source interfaces and source VSANs before the upgrade. After the upgrade, the source VSANs were removed (rule 1).

- If interface level VSAN filters are configured in source interfaces, then the source interfaces are also removed from the session. If this interface is configured in both directions, it will be removed from both directions. For example,

Before Release 1.0(4):

```
Session 2 (active)
  Destination is fc1/9
  No session filters configured
  Ingress (rx) sources are
    vsans 12
    fc1/6 (vsan 1-20),
  Egress (tx) sources are
    fc1/6 (vsan 1-20),
```

Once upgraded to Release 1.1(1):

```
Session 2 (inactive as no active sources)
  Destination is fc1/9
  No session filters configured
  No ingress (rx) sources
  No egress (tx) sources
```

Session 2 had a source VSAN 12 and a source interface fc1/6 with VSAN filters specified in Release 1.0(4). When upgraded to Release 1.1(1) the following changes are made:

- The source VSAN (VSAN 12) is removed (rule 1)
- The source interface fc1/6 had VSAN filters specified—it is also removed (rule 2).

**Note**

The deprecated configurations are removed from persistent memory, once a switchover or a new startup configuration is implemented.

Monitoring Traffic Using Fibre Channel Analyzers

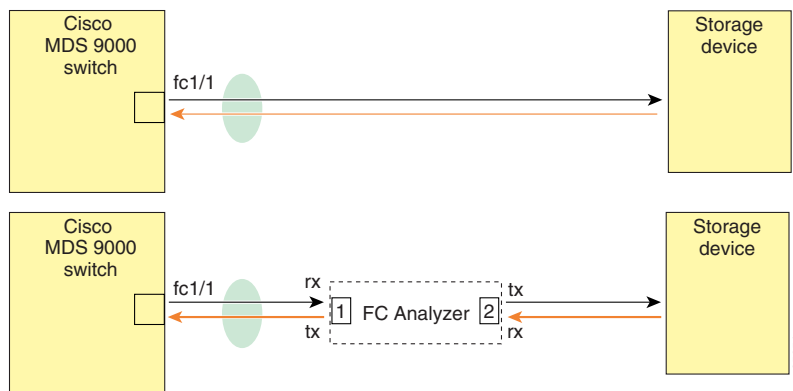
You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is specially useful in troubleshooting scenarios when traffic disruption changes the problem environment and makes it difficult to reproduce the problem.

Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 as shown in [Figure 28-5](#).

Figure 28-5 Fibre Channel Analyzer Usage Without SPAN

FC Analyzer usage without SPAN



This type of connection has the following limitations:

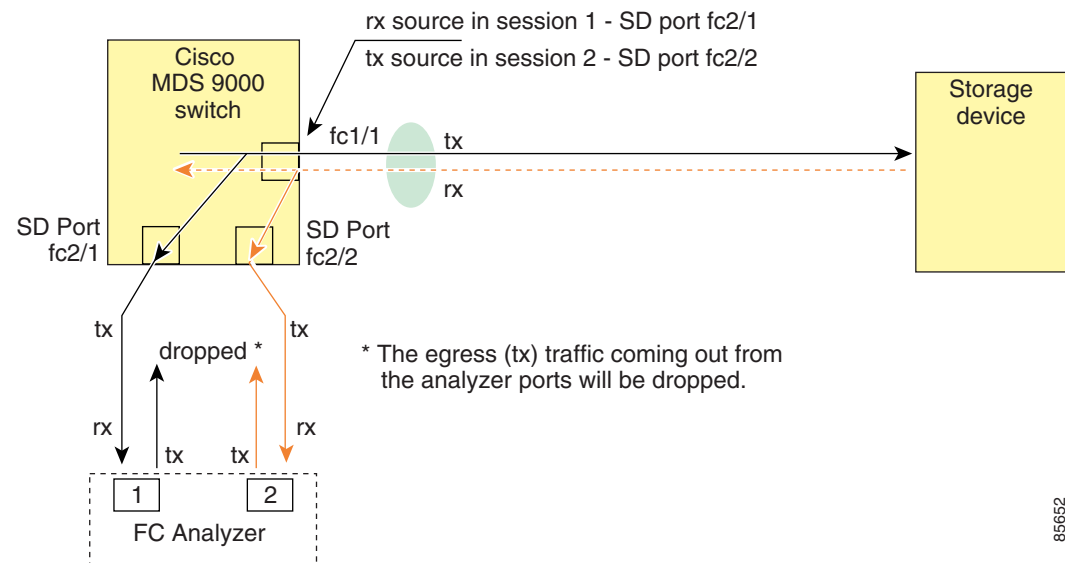
- Requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

Using SPAN

Using SPAN you can capture the same traffic scenario shown in [Figure 28-5](#) without any traffic disruption. The Fibre Channel analyzer uses the ingress (rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2, to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer as shown in [Figure 28-6](#).

Figure 28-6 Fibre Channel Analyzer Using SPAN



Configuring Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in [Figure 28-6](#), follow these steps:

- | | |
|---------------|---|
| Step 1 | Configure SPAN on interface fc1/1 in the ingress (rx) direction to send traffic on SD port fc2/1 using session 1. |
| Step 2 | Configure SPAN on interface fc1/1 in the egress (tx) direction to send traffic on SD port fc2/2 using session 2. |
| Step 3 | Physically connect fc2/1 to port 1 on the Fibre Channel analyzer. |
| Step 4 | Physically connect fc2/2 to port 2 on the Fibre Channel analyzer. |

To configure SPAN on the source and destination interfaces, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# span session 1 switch(config-span)#	Creates the SPAN session 1.

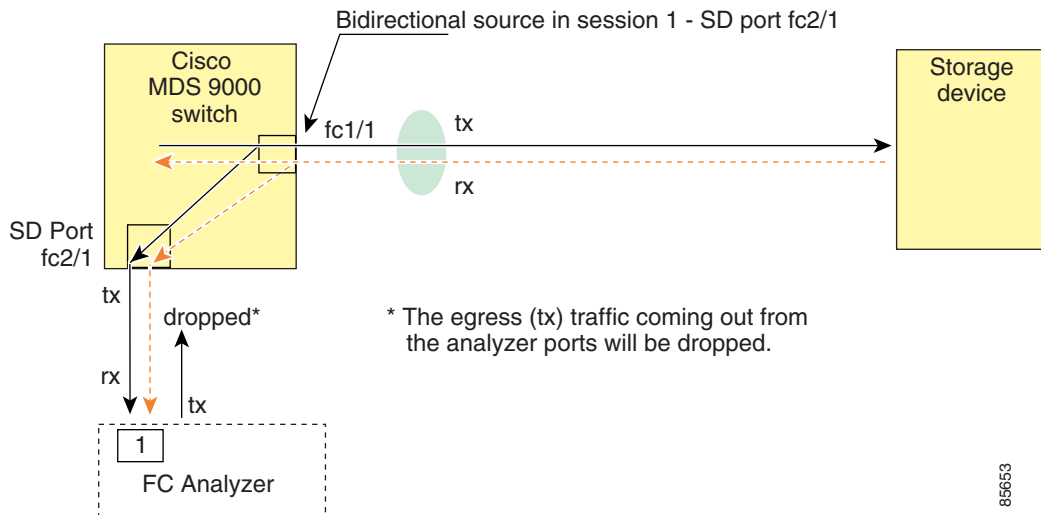
	Command	Purpose
Step 3	switch(config-span)## destination interface fc2/1	Configures the destination interface fc2/1.
Step 4	switch(config-span)# source interface fc1/1 rx	Configures the source interface fc1/1 in the ingress direction.
Step 5	switch(config)# span session 2 switch(config-span)#	Creates the SPAN session 2.
Step 6	switch(config-span)## destination interface fc2/2	Configures the destination interface fc2/2.
Step 7	switch(config-span)# source interface fc1/1 tx	Configures the source interface fc1/1 in the egress direction.

Using a Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface as shown in Figure 28-6. You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

Figure 28-7 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress direction. This setup is more advantageous and cost-effective than the setup shown in Figure 28-6—it uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 28-7 Fibre Channel Analyzer Using a Single SD Port



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

To configure SPAN on a single SD port, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# span session 1 switch(config-span)#	Creates the SPAN session 1.

	Command	Purpose
Step 3	switch(config-span)## destination interface fc2/1	Configures the destination interface fc2/1.
Step 4	switch(config-span)# source interface fc1/1	Configures the source interface fc1/1 on the same SD port.

Displaying SPAN Information

Use the **show span** command to display configured SPAN information. See Examples 28-1 to 28-4.

Example 28-1 Displays SPAN Sessions in a Brief Format

```
switch# show span session brief
```

Session	Admin State	Oper State	Destination Interface
7	no suspend	active	fc2/7
1	suspend	inactive	not configured
2	no suspend	inactive	fc3/1

Example 28-2 Displays a Specific SPAN Session Details

```
switch# show span session 7
Session 7 (active)
  Destination is fc2/7
  No session filters configured
  No ingress (rx) sources
  Egress (tx) sources are
    port-channel 7,
```

Example 28-3 Displays ALL SPAN Sessions

```
switch# show span session
Session 1 (inactive as no destination)
Destination is not specified
  Session filter vsans are 1
  No ingress (rx) sources
  No egress (tx) sources
Session 2 (active)
  Destination is fc9/5
  No session filters configured
  Ingress (rx) sources are
    vsans 1
  No egress (tx) sources
Session 3 (admin suspended)
  Destination is not configured
  Session filter vsans are 1-20
  Ingress (rx) sources are
    fc3/2, fc3/3, fc3/4, fcip 51,
    port-channel 2, sup-fc0,
  Egress (tx) sources are
    fc3/2, fc3/3, fc3/4, sup-fc0,
```

Example 28-4 Displays an SD-port Interface with Encapsulation Enabled

```
switch# show int fc9/32
fc9/32 is up
  Hardware is Fibre Channel
  Port WWN is 22:20:00:05:30:00:49:5e
  Admin port mode is SD
  Port mode is SD
  Port vsan is 1
  Speed is 1 Gbps
  Receive Buffer Size is 2112
  Encapsulation is eisl <----- Displays the enabled encapsulation status
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes, 0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```

Default SPAN Settings

Table 28-1 lists the default settings for SPAN parameters

Table 28-1 Default SPAN Configuration Parameters

Parameters	Default
SPAN session	Active.
If filters are not specified	SPAN traffic includes traffic through a specific interface from all active VSANs.
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.

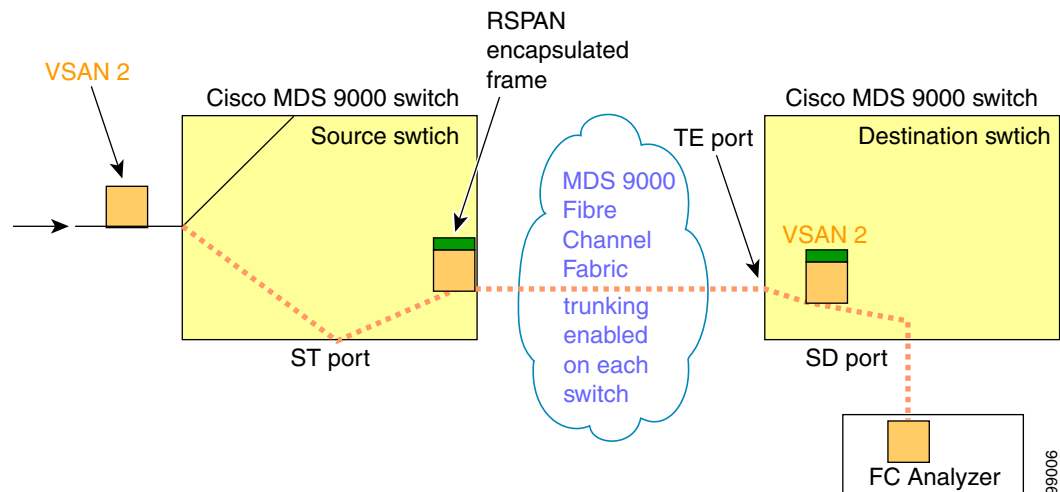
Remote SPAN

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for any SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types (see [Figure 28-8](#)):

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 28-8 RSPAN Transmission



Advantages to Using RSPAN

The RSPAN features has the following advantages:

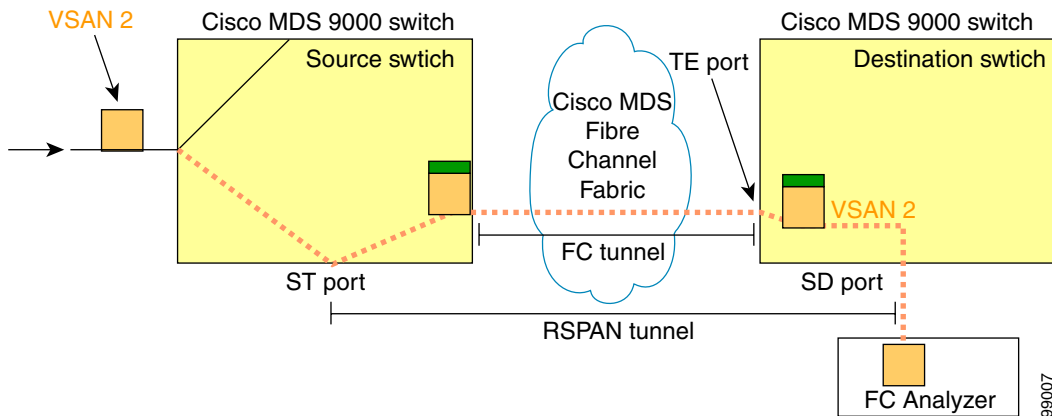
- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost-effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

FC and RSPAN Tunnels

A FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to a ST port in the source switch and map the same FC tunnel to a SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as a RSPAN tunnel (see [Figure 28-9](#)).

Figure 28-9 FC and RSPAN Tunnel



Guidelines to Configure RSPAN

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it will be dropped.
- The FC tunnel's IP address must reside in the same subnet as the VSAN interface (see [Chapter 20, "Configuring IP Services"](#)).
- The following configurations must be performed on *each* switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented
 - Trunking must be enabled (the **trunk protocol enable** command is enabled by default).
 - VSAN interface must be configured (the **interface vsan** command).
 - The Fibre Channel tunnel feature must be enabled (the **fc-tunnel enable** command is disabled by default).
 - IP routing must be enabled (the **ip routing** command is disabled by default).



Note

If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.

- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.

ST Port Characteristics

ST port have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.
- A ST port can only be bound to one FC tunnel.
- ST ports cannot be used for any other purpose other than to carry RSPAN traffic.
- ST Ports cannot be configured using Advanced Services Modules (ASMs).

Configuring RSPAN

The RSPAN tunnel begins in the source switch and terminates in the destination switch. This section assumes Switch S to be the source and Switch D to be the destination.

**Note**

Besides the source and destination switches, the VSAN must also be configured in each MDS switch in the Fibre Channel fabric, if they exist.

To monitor network traffic using the RSPAN feature, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Create VSAN interfaces in destination switch (Switch D) and source switch (Switch S) to facilitate the Fibre Channel tunnel (FC tunnel) creation. |
| Step 2 | Enable the FC tunnel in each switch in the end-to-end path of the tunnel. |
| Step 3 | Initiate the FC tunnel (in Switch S) and map the tunnel to the VSAN interface's IP address (in Switch D) so all RSPAN traffic from the tunnel is directed to the SD port. |
| Step 4 | Configure SD ports for SPAN monitoring in the destination switch (Switch D). |
| Step 5 | Configure the ST port in the source switch (Switch S) and bind the ST port to the FC tunnel. |
| Step 6 | Create a RSPAN session in the source switch (in Switch S) to monitor network traffic. |
-

Configuration in the Source Switch

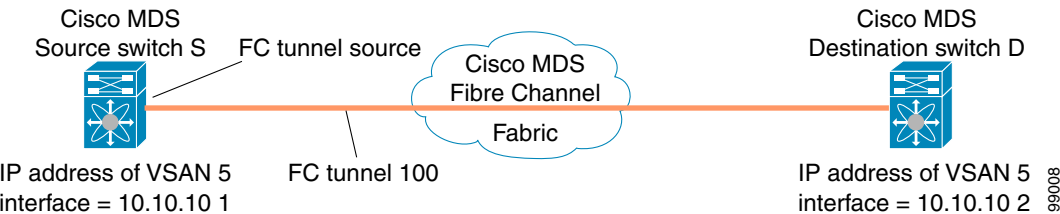
This section identifies the tasks that must be performed in the source switch (Switch D):

- [Creating VSAN Interfaces](#)
- [Enabling FC Tunnels](#)
- [Initiating the FC Tunnel](#)
- [Configuring the ST Port](#)
- [Configure a RSPAN Session](#)

Creating VSAN Interfaces

Figure 28-10 depicts a basic FC tunnel configuration.

Figure 28-10 FC Tunnel Configuration




Note

This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interfaces in the source switch for the scenario in Figure 28-10, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# interface vsan 5 switchS(config-if)#	Configures the specified VSAN interface (VSAN 5) in the source switch (switch S).
Step 3	switchS(config-if)# ip address 10.10.10.1 255.255.255.0	Configures the IP address and subnet for the VSAN interface 5 in the source switch (switch S).
Step 4	switchS(config-if)# no shutdown	Enables traffic flow through this interface.

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# fc-tunnel enable	Enables the FC tunnel feature (disabled by default).


Note

Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Initiating the FC Tunnel

To initiate the FC tunnel in the source switch for the scenario in Figure 28-10, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# interface fc-tunnel 100 switchS(config-if)#	Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255.

	Command	Purpose
Step 3	switchS(config-if)# source 10.10.10.1	Maps the IP address of the source switch (switch S) to the FC tunnel (100).
Step 4	switchS(config-if)# destination 10.10.10.2	Maps the IP address of the destination switch (switch D) to the FC tunnel (100).
Step 5	switchS(config-if)# no shutdown	Enables traffic flow through this interface.

**Tip**

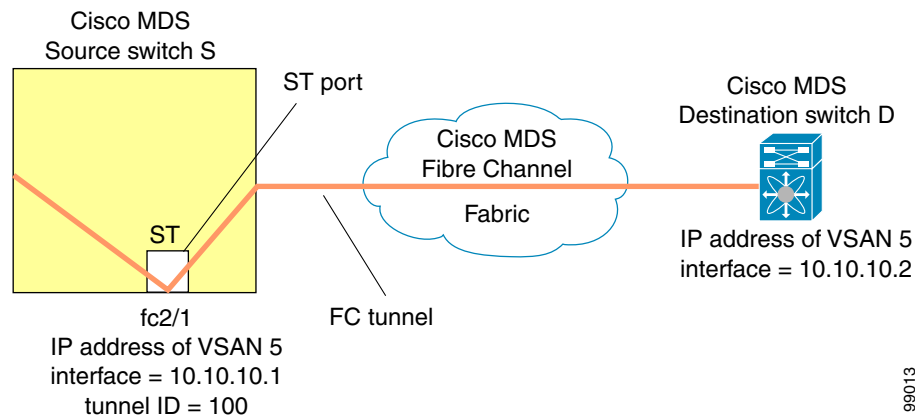
The interface will not be operationally up until the FC tunnel mapping is configured in the destination switch.

Configuring the ST Port

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes a RSPAN tunnel once the binding and mapping is complete.

Figure 28-11 depicts a basic FC tunnel configuration.

Figure 28-11 Binding the FC Tunnel



To configure an ST port for the scenario in Figure 28-11, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# interface fc2/1	Configures the specified interface.
Step 3	switchS(config-if)# switchport mode ST	Configures the ST port-mode for interface fc2/1.
Step 4	switchS(config-if)# switchport speed 2000	Configures the ST port speed to 2000 Mbps.
Step 5	switchS(config-if)# rspan-tunnel interface fc-tunnel 100	Associates and binds the ST port with the RSPAN tunnel (100).
Step 6	switchS(config-if)# no shutdown	Enables traffic flow through this interface.

**Note**

ST ports cannot be configured using Advanced Services Modules (ASMs).

Configure a RSPAN Session

A RSPAN session is similar to a SPAN session, with the destination interface being a RSPAN tunnel. To configure a RSPAN session in the source switch for the scenario in [Figure 28-11](#), follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# span session 2 switchS(config-span)#	Configures the specified SPAN session (2). If the session does not exist, it will be created. The session ID ranges from 1 to 16.
Step 3	switchS(config-span)# destination interface fc-tunnel 100	Configures the specified RSPAN tunnel (100) in a session.
Step 4	switch(config-span)# source interface fc1/1	Configures the source interface (fc1/1) for this session and spans the traffic from interface fc1/1 to RSPAN tunnel 100.

Configuration in All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel:

- [Configuring VSAN Interfaces](#)
- [Enabling FC Tunnels](#)
- [Enabling IP Routing](#)

Configuring VSAN Interfaces

[Figure 28-13](#) depicts a RSPAN tunnel configuration terminating in the destination switch (Switch D).



Note

This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in [Figure 28-13](#), follow these steps:

	Command	Purpose
Step 1	switchD# config t	Enters configuration mode.
Step 2	switchD(config)# interface vsan 5 switchD(config-if)#	Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D).
Step 3	switchD(config-if)# ip address 10.10.10.2 255.255.255.0	Configures the IP address and subnet for the VSAN interface in the destination switch (Switch D).
Step 4	switch(config-if)# no shutdown	Enables traffic flow to administratively allow traffic (provided the operational state is up).

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# fc-tunnel enable	Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs ranges from 1 to 255.



Note

Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric (see [“Enabling IP Routing” section on page 20-11](#)). This step is required to setup the FC tunnel.

Configuration in the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D):

- [Configuring VSAN Interfaces](#)
- [Configuring the SD Port](#)
- [Mapping the FC Tunnel](#)

Configuring VSAN Interfaces

[Figure 28-13](#) depicts a RSPAN tunnel configuration terminating in the destination switch (Switch D).



Note

This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in [Figure 28-13](#), follow these steps:

	Command	Purpose
Step 1	switchD# config t	Enters configuration mode.
Step 2	switchD(config)# interface vsan 5 switchD(config-if)#	Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D).
Step 3	switchD(config-if)# ip address 10.10.10.2 255.255.255.0	Configures the IP address and subnet for the VSAN interface in the destination switch (Switch D).
Step 4	switchD(config-if)# no shutdown	Enables traffic flow to administratively allow traffic (provided the operational state is up).

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# fc-tunnel enable	Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255.



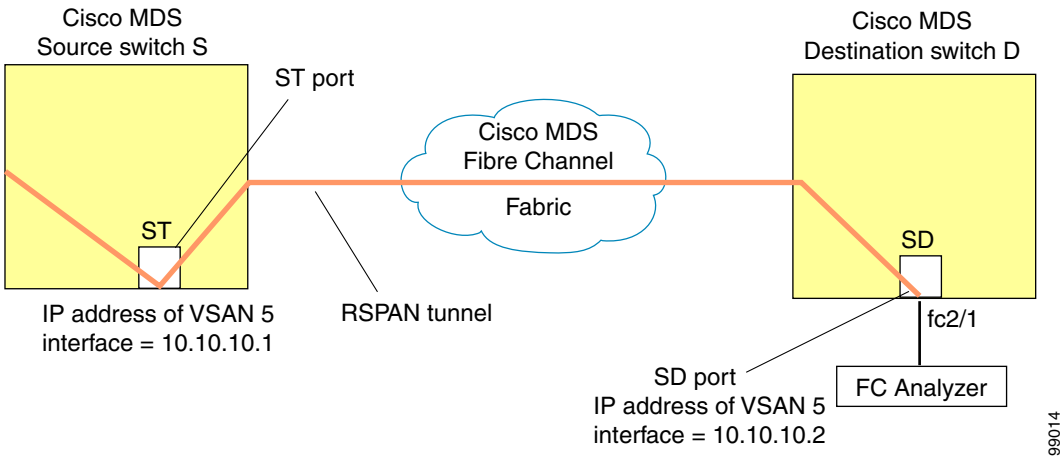
Note

Be sure to enable this feature in each switch in the end-to-end path in the tunnel.

Configuring the SD Port

The SD port in the destination switch enables the FC Analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. Figure 28-12 depicts a RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 28-12 RSPAN Tunnel Configuration



To configure an SD port for the scenario in Figure 28-12, follow these steps:

	Command	Purpose
Step 1	switchD# config t	Enters configuration mode.
Step 2	switchD(config)# interface fc2/1	Configures the specified interface.
Step 3	switchD(config-if)# switchport mode SD	Configures the SD port-mode for interface fc2/1.
Step 4	switchD(config-if)# switchport speed 2000	Configures the SD port speed to 2000 Mbps.
Step 5	switchD(config-if)# no shutdown	Enables traffic flow through this interface.



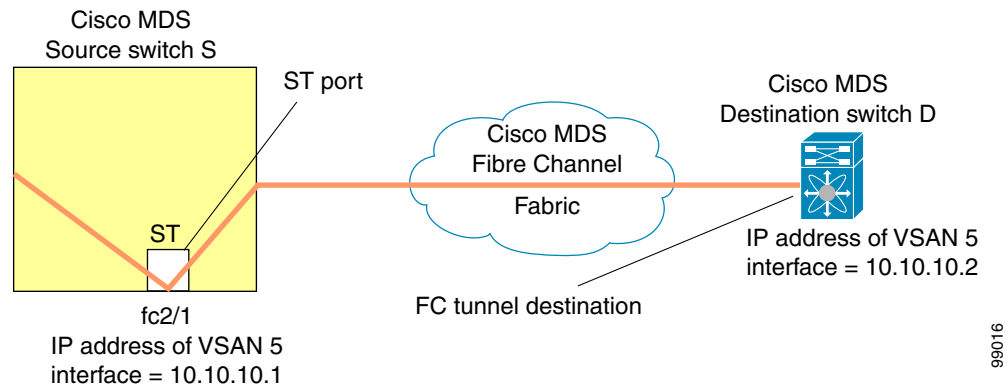
Note

SD ports cannot be configured using Advanced Services Modules (ASMs).

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch (see [Figure 28-13](#)).

Figure 28-13 FC Tunnel Configuration



To terminate the FC tunnel in the destination switch for the scenario in [Figure 28-13](#), follow these steps:

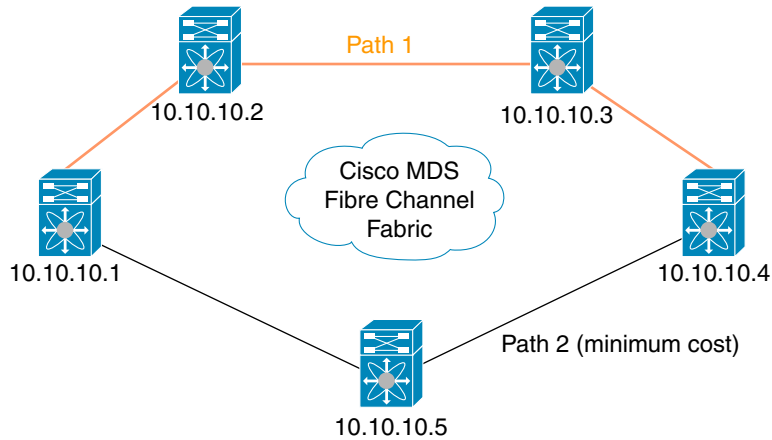
	Command	Purpose
Step 1	switchD# <code>conf t</code>	Enters configuration mode.
Step 2	switchD(config)# <code>fc-tunnel</code> switchD(config)# <code>tunnel-id-map 100 interface fc2/1</code>	Terminates the FC tunnel (100) in the destination switch (switch D). The tunnel ID range is from 1 to 255.

Configuring An Explicit Path

You can specify an explicit path through the Cisco MDS Fibre channel fabric (source-based routing), use the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the fc-tunnel to always take one path to the destination switch. The software then use this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths available. In a RSPAN situation, you can specify the explicit-path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch (see [Figure 28-14](#)).

Figure 28-14 Explicit Path Configuration



The explicit path must be created in the source switch. To configure an explicit path, you must first create the path and then configure the use of any one path. If an explicit path is not configured, the minimum cost path is used by default. If an explicit path is configured and is functioning, the specified path is used.

To create an explicit path for the scenario in Figure 28-14, follow these steps:

	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# fc-tunnel explicit-path Path1 switch(config-explicit-path)#	Places you at the explicit path prompt for the path named Path 1.
Step 3	switchS(config-explicit-path)# next-address 10.10.10.2 strict switchS(config-explicit-path)# next-address 10.10.10.3 strict switchS(config-explicit-path)# next-address 10.10.10.4 strict	Specifies that the next hop VSAN interface IP addresses and the previous hops specified in the explicit path do not require direct connection.
Step 4	switchS(config)# fc-tunnel explicit-path Path2 switch(config-explicit-path)#	Places you at the explicit path prompt for Path2.
Step 5	switchS(config-explicit-path)# next-address 10.10.10.5 strict switchS(config-explicit-path)# next-address 10.10.10.4 strict	Specifies that the next hop VSAN interface IP addresses and the previous hops specified in the explicit path does not require direct connection.
Step 6	switchS(config)# fc-tunnel explicit-path Path3 switch(config-explicit-path)#	Places you at the explicit path prompt for Path3.
Step 7	switchS(config-explicit-path)# next-address 10.10.10.3 loose	Configures a minimum cost path in which the 10.10.10.3 IP address exists. Note In Figure 28-14, Path3 is the same as Path1—10.10.10.3 exists in Path 1. Using the loose option, you can achieve the same effect with one command instead of issuing three commands (using the strict option) in Step 3.

To reference the explicit path, follow these steps:

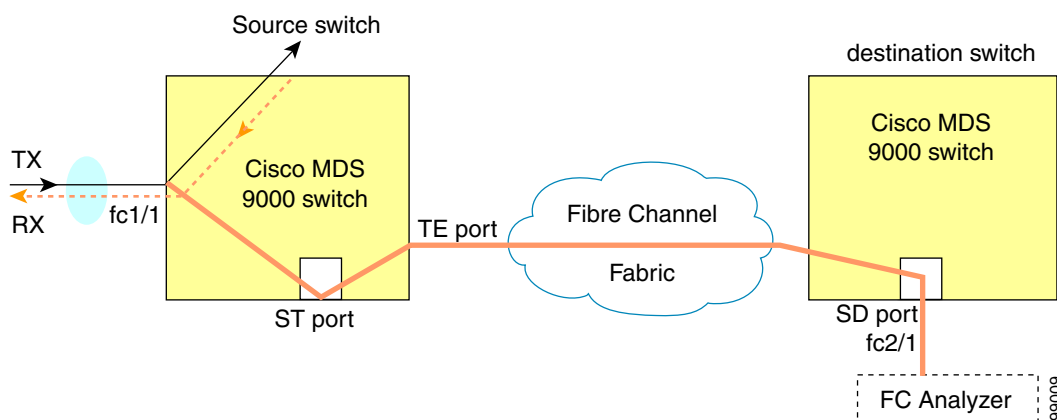
	Command	Purpose
Step 1	switchS# config t	Enters configuration mode.
Step 2	switchS(config)# interface fc-tunnel 100	References the tunnel ID for Path1.
Step 3	switchS(config)# explicit-path Path1	Links Path1 to the tunnel ID.

This configuration explicitly specifies Path 1 to be used for the RSPAN traffic. Refer to RFC 3209 for further details on explicit paths and source based routing.

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. Figure 28-7 shows a RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress direction.

Figure 28-15 Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Sample Scenarios



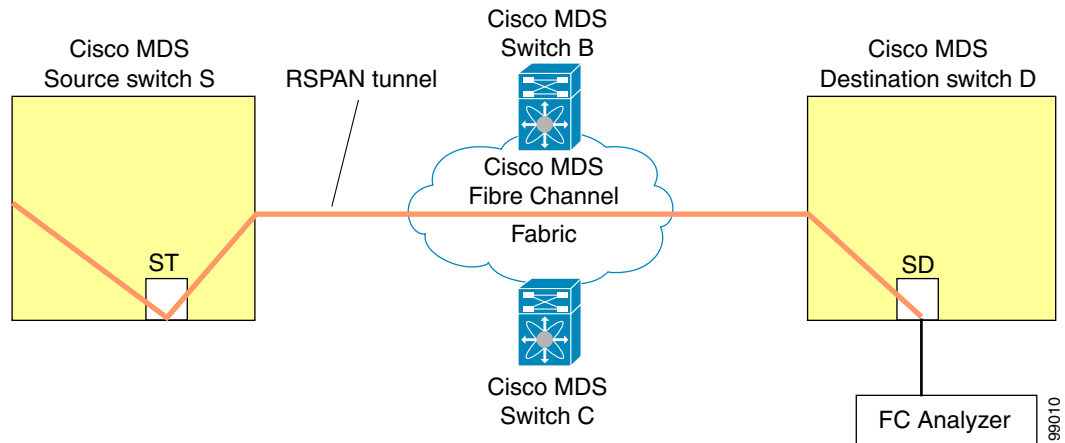
Note

RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. A RSPAN tunnel is configured as a destination interface for SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel (see [Figure 28-16](#)).

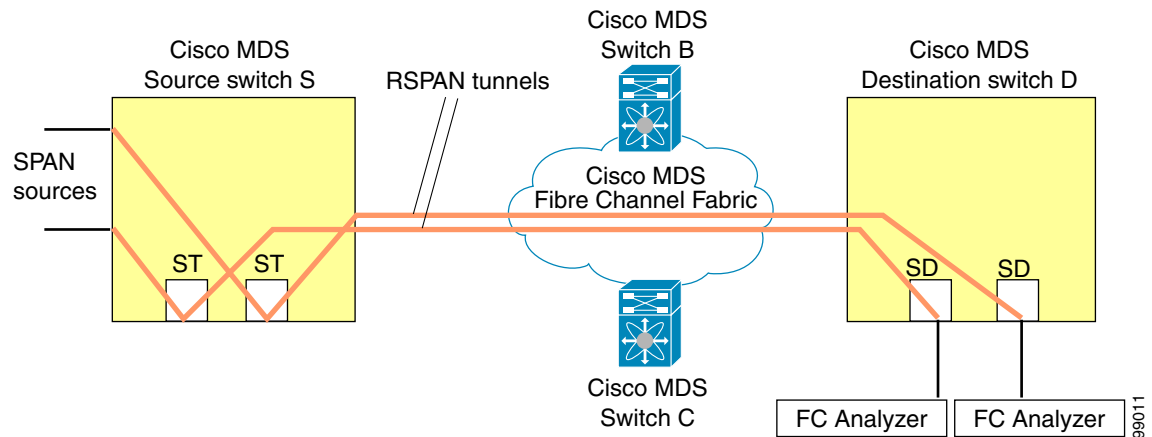
Figure 28-16 RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel



Single Source with Multiple RSPAN Tunnels

[Figure 28-17](#) displays two separate RSPAN tunnels configured between Switches S and N. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for trouble shooting purposes.

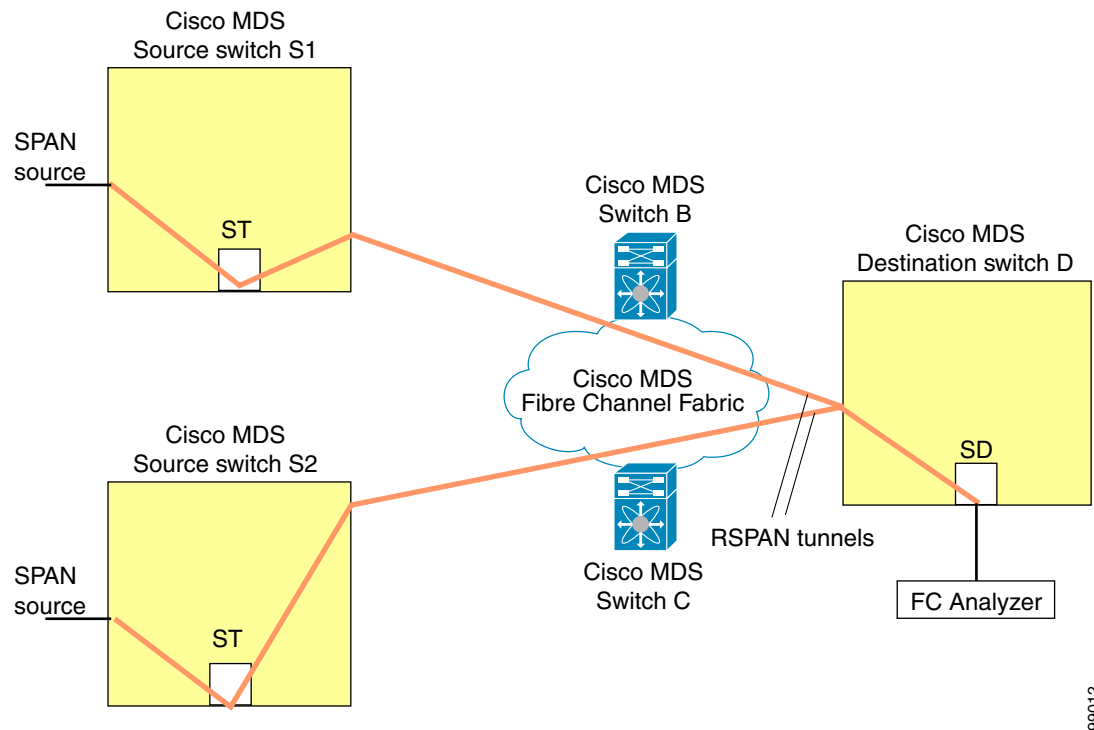
Figure 28-17 RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels



Multiple Sources with Multiple RSPAN Tunnels

Figure 28-18 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 28-18 RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels



This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

Displaying RSPAN Information

Use the **show** commands to display configured RSPAN information. See Examples 28-1 to 28-4.

Example 28-5 Displays ST Port Interface Information

```
switch# show interface brief
```

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	Oper Mode	Oper Speed (Gbps)	Port-channel
fc1/1	1	auto	on	trunking	TE	2	--
...							
fc1/14	1	auto	on	trunking	TE	2	--
fc1/15	1	ST	on	up	ST	2	--
...							
fc2/9	1	auto	on	trunking	TE	2	port-channel 21

```

fc2/10      1      auto  on      trunking  TE      2      port-channel 21
...
fc2/13      999    auto  on      up         F       1      --
fc2/14      999    auto  on      up         FL      1      --
fc2/15      1      SD    --      up         SD      2      --
fc2/16      1      auto  on      trunking  TE      2      --
-----
Interface          Status                      Speed
                        (Gbps)
-----
sup-fc0            up                          1
-----
Interface          Status      IP Address          Speed      MTU
-----
mgmt0              up          172.22.36.175/22    100 Mbps   1500
-----
Interface          Status      IP Address          Speed      MTU--
-----
vsan5              up          10.10.10.1/24       1 Gbps     1500
-----
Interface          Vsan        Admin              Status      Oper      Oper
                        Trunk              Mode        Mode        Speed
                        Mode              (Gbps)
-----
port-channel 21    1           on                 trunking    TE        4
-----
Interface          Status      Dest IP Addr      Src IP Addr      TID      Explicit Path
-----
fc-tunnel 100      up          10.10.10.2        10.10.10.1      100

```

Example 28-6 Displays Detailed Information for the ST Port Interface

```

switch# show interface fc1/11
fc1/11 is up
  Hardware is Fibre Channel
  Port WWN is 20:0b:00:05:30:00:59:de
  Admin port mode is ST
  Port mode is ST
  Port vsan is 1
  Speed is 1 Gbps
  Rspan tunnel is fc-tunnel 100
  Beacon is turned off
  5 minutes input rate 248 bits/sec, 31 bytes/sec, 0 frames/sec
  5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
    6862 frames input, 444232 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
      0 too long, 0 too short
    6862 frames output, 307072 bytes
      0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits

```

Example 28-7 Displays the FC Tunnel Status

```

switch# show fc-tunnel
fc-tunnel is enabled

```

Example 28-8 Displays FC Tunnel Egress Mapping Information

```
switch# show fc-tunnel tunnel-id-map
tunnel id egress interface
      150      fc3/1
      100      fc3/1
```

**Note**

Multiple tunnel IDs can terminate at the same interface.

Example 28-9 Displays FC Tunnel Explicit Mapping Information

```
switch# show fc-tunnel explicit-path
Explicit path name: Alternatel
      10.20.1.2 loose
      10.20.1.3 strict
Explicit path name: User2
      10.20.50.1 strict
      10.20.50.4 loose
```

Example 28-10 Displays SPAN Mapping Information

```
switch# show span session
Session 2 (active)
  Destination is fc-tunnel 100
  No session filters configured
  Ingress (rx) sources are
    fc2/16,
  Egress (tx) sources are
    fc2/16,
```

Example 28-11 Displays the FC Tunnel Interface

```
switch# show interface fc-tunnel 200
fc-tunnel 200 is up
Dest   IP Addr: 200.200.200.7   Tunnel ID: 200
Source IP Addr: 200.200.200.4   LSP ID: 1
Explicit Path Name:
```




Advanced Features and Concepts

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [Configuring FC Timers, page 29-2](#)
- [Invoking the fctrace Feature, page 29-4](#)
- [Invoking the fcping Feature, page 29-5](#)
- [Configuring a Fabric Analyzer, page 29-7](#)
- [Configuring World Wide Names, page 29-18](#)
- [Allocating Flat FC IDs, page 29-19](#)
- [Enabling Loop Monitoring, page 29-20](#)
- [Configuring the Switch for Interoperability, page 29-21](#)
- [Using the show tech-support Command, page 29-27](#)

Configuring FC Timers

The **ftimer** command modifies Fibre Channel protocol related timer values for the switch.

You can use the **ftimer** command in configuration mode to configure the following TOVs:

- Distributed services TOV (D_S_TOV)—the valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—the valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—the valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



Note

The Fabric stability TOV (F_S_TOV) constant cannot be configured.

Configuring Timers Across All VSANs

To configure FC timers across all VSANs, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)	Enters configuration mode.
Step 2	switch(config)# ftimer R_A_TOV 6000	Configures the R_A_TOV value for all VSANs to be 6000 ms. This type of configuration is not permitted unless all VSANs are suspended



Caution

The D_S_TOV, E_D_TOV, and R_A_TOV values cannot be globally changed unless all VSANs in the switch are suspended.



Note

If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

Configuring Timers Per-VSAN

You can also issue the **ftimer** command for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended or activated when their timer values are changed.



Caution

You cannot perform a nondisruptive downgrade to any earlier version which does not support per-VSAN FC timers.

To configure per-VSAN FC timers, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)	Enters configuration mode.
Step 2	switch(config)# fctimer D_S_TOV 6000 vsan 2 Warning: The vsan will be temporarily suspended when updating the timer value This configuration would impact whole fabric. Do you want to continue? (y/n) y Since this configuration is not propagated to other switches, please configure the same value in all the switches	Configures the D_S_TOV value to be 6000 ms for VSAN 2. Suspends the VSAN temporarily. You have the option to abort this command, if required.

**Note**

This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to SAN-OS Release 1.2(x) or 1.1(x) after the timer is configured for a VSAN, an error message is issued to warn the user about strict incompatibilities. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide* for further information.

Displaying Configured FC Timer Values

Use the **show fctimer** command to display the configured FC timer values (see Examples 29-1 and 29-2).

Example 29-1 Displays Configured Global TOVs

```
switch# show fctimer
F_S_TOV   D_S_TOV   E_D_TOV   R_A_TOV
-----
5000 ms   5000 ms   2000 ms   10000 ms
```

**Note**

The F_S_TOV constant, though not configured, is displayed in the output of the **show fctimer** command.

Example 29-2 Displays Configured TOVs for a Specified VSAN

```
switch# show fctimer vsan 10
vsan no.  F_S_TOV   D_S_TOV   E_D_TOV   R_A_TOV
-----
10         5000 ms   5000 ms   3000 ms   10000 ms
```

Invoking the fctrace Feature

The fctrace feature allows you to:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

You can invoke fctrace by providing the FC ID, the N port, or the NL port WWN of the destination. The frames are routed normally as long as they are forwarded through TE ports.

Once the frame reaches the edge of the fabric (the F port or FL port connected to the end node with the given port WWN or the FC ID), the frame is looped back (swapping the source ID and the destination ID) to the originator.

If the destination cannot be reached the path discovery starts, which traces the path up to the point of failure.



Note

The fctrace feature works only on TE Ports. Make sure that only TE ports exist in the path to the destination. In case there is an E Port in the path, the fctrace frame is dropped by that switch. Also, fctrace times out in the originator, and path discovery does not start.

To perform a fctrace operation, follow this step:

	Command	Purpose
Step 1	<pre>switch# fctrace fcid 0xd70000 vsan 1 Route present for : 0xd70000 20:00:00:0b:46:00:02:82(0xffffcd5) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7)</pre>	Invokes fctrace for the specified FC ID of the destination N port
	<pre>switch# fctrace pwn 21:00:00:e0:8b:06:d9:1d vsan 1 timeout 5 Route present for : 21:00:00:e0:8b:06:d9:1d 20:00:00:0b:46:00:02:82(0xffffcd5) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7)</pre>	<p>Invokes fctrace using the pWWN of the destination N port</p> <p>By default the period to wait before timing out is 5 seconds, The range is from one through 10 seconds.</p>



Note

You cannot use the fctrace feature in a locally configured VSAN interface (IPFC interface), but you can trace the route to a VSAN interface configured in other switches.

Invoking the fcping Feature

The fcping feature verifies reachability of a node by checking its end-to-end connectivity. You can invoke the fcping feature by providing the FC ID or the destination port WWN information.

To perform a fcping operation, follow these steps:

	Command	Purpose
Step 1	<pre>switch# fcping fcid 0xd70000 vsan 1 28 bytes from 0xd70000 time = 730 usec 28 bytes from 0xd70000 time = 165 usec 28 bytes from 0xd70000 time = 262 usec 28 bytes from 0xd70000 time = 219 usec 28 bytes from 0xd70000 time = 228 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 165/270/730 usec</pre>	Performs a fcping operation for the specified pWWN or the FCID of the destination. By default, five frames are sent.
	<pre>switch# fcping fcid 0xd70000 vsan 1 count 10 28 bytes from 0xd70000 time = 730 usec 28 bytes from 0xd70000 time = 165 usec 28 bytes from 0xd70000 time = 262 usec 28 bytes from 0xd70000 time = 219 usec 28 bytes from 0xd70000 time = 228 usec 28 bytes from 0xd70000 time = 230 usec 28 bytes from 0xd70000 time = 230 usec 28 bytes from 0xd70000 time = 225 usec 28 bytes from 0xd70000 time = 229 usec 28 bytes from 0xd70000 time = 183 usec 10 frames sent, 10 frames received, 0 timeouts Round-trip min/avg/max = 165/270/730 usec</pre>	Sets the number of frames to be sent using the count option. The range is from 0 through 2147483647. A value of 0 will ping forever.
	<pre>switch# fcping fcid 0xd500b4 vsan 1 timeout 10 28 bytes from 0xd500b4 time = 1345 usec 28 bytes from 0xd500b4 time = 417 usec 28 bytes from 0xd500b4 time = 340 usec 28 bytes from 0xd500b4 time = 451 usec 28 bytes from 0xd500b4 time = 356 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 340/581/1345 usec</pre>	Sets the timeout value. The default period to wait is 5 seconds. The range is from 1 through 10 seconds.
Step 2	<pre>switch# fcping fcid 0x010203 vsan 1 No response from the N port. switch# fcping pwn 21:00:00:20:37:6f:db:dd vsan 1 28 bytes from 21:00:00:20:37:6f:db:dd time = 1454 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 471 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 372 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 364 usec 28 bytes from 21:00:00:20:37:6f:db:dd time = 1261 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 364/784/1454 usec</pre>	<p>Issues a No response from the N port message even when the N port or NL port is active. This is due to resource exhaustion at the N port or NL port.</p> <p>Retry the command a few seconds later.</p>

Verifying Switch Connectivity

You can also use the **fcping fcid** command to verify connectivity to a destination switch.


Note

The FC ID variable used in this procedure is the domain controller address; it is not a duplication of the domain ID.

To verify connectivity to a destination switch, follow these steps:

	Command	Purpose
Step 1	<pre>switch# show fcdomain domain-list vsan 200 Number of domains: 7 Domain ID WWN ----- 0x01(1) 20:c8:00:05:30:00:59:df [Principal] 0x02(2) 20:c8:00:0b:5f:d5:9f:c1 0x6f(111) 20:c8:00:05:30:00:60:df 0xda(218) 20:c8:00:05:30:00:87:9f [Local] 0x06(6) 20:c8:00:0b:46:79:f2:41 0x04(4) 20:c8:00:05:30:00:86:5f 0x6a(106) 20:c8:00:05:30:00:f8:e3</pre>	Displays the destination switch’s domain ID. To obtain the domain controller address, concatenate the domain ID with FFFC . For example, if the domain ID is 0xda(218) , the concatenated ID is 0xffffcda .
Step 2	<pre>switch# fcping fcid 0xFFFFCDA vsan 200 28 bytes from 0xFFFFCDA time = 298 usec 28 bytes from 0xFFFFCDA time = 260 usec 28 bytes from 0xFFFFCDA time = 298 usec 28 bytes from 0xFFFFCDA time = 294 usec 28 bytes from 0xFFFFCDA time = 292 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 260/288/298 usec</pre>	Verifies reachability of the destination switch by checking its end-to-end connectivity.

Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. While existing Fibre Channel analyzers can capture traffic at wire rate speed. They are expensive and support limited frame decoding. Also, to snoop traffic, the existing analyzers disrupt the traffic on the link while the analyzer is inserted into the link.

Cisco has brought protocol analysis within a storage network to a new capability level with the Cisco Fabric Analyzer. You can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be local to the point of analysis.

Cisco's Fibre Channel protocol analyzer is based on two popular public-domain software applications:

- libpcap—You can obtain more information from <http://www.tcpdump.org>.
- Ethereal—You can obtain more information from <http://www.ethereal.com>.

**Note**

Cisco's Fabric Analyzer is useful in capturing and decoding control traffic, not data traffic. It is suitable for control path captures, and is not intended for high-speed data path captures.

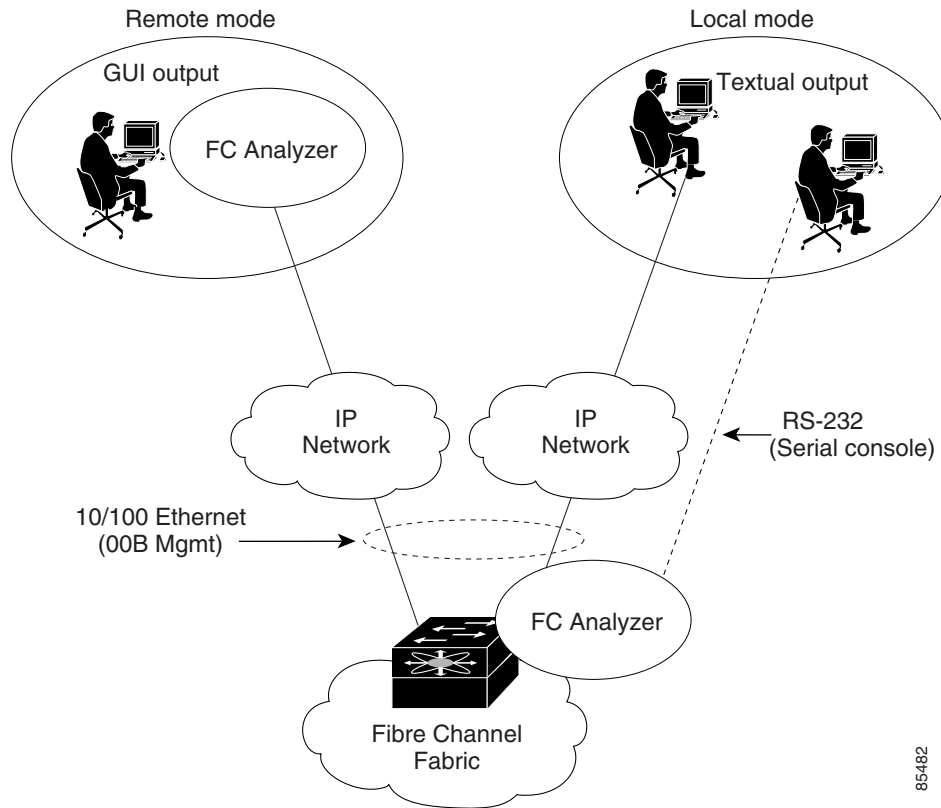
This section explains the following topics:

- [About the Cisco Fabric Analyzer, page 29-7](#)
- [Configuring the Cisco Fabric Analyzer, page 29-9](#)
- [Viewing Display Filters Information, page 29-12](#)
- [Clearing Configured fcanalyzer Information, page 29-11](#)
- [Display Filters, page 29-12](#)

About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer comprises two separate components (see [Figure 29-1](#)):

- A software that runs on the Cisco MDS 9000 Family switch and supports two modes of capture:
 - a text-based analyzer that supports local capture and decodes captured frames
 - a daemon that supports remote capture
- A GUI-based client that runs on a host that supports libpcap such as Windows or Linux and communicates with the remote capture daemon in a Cisco MDS 9000 Family switch.

Figure 29-1 Cisco Fabric Analyzer Usage

Local Text-Based Capture

This component is a command-line driven text-based interface that captures traffic to and from the supervisor module in a Cisco MDS 9000 switch. It is a fully-functional decoder that is useful for quick debug purposes or for use when the remote capture daemon is not enabled. Additionally, because this tool is accessed from within the Cisco MDS 9000 switch, it is protected by the roles-based policy that limits access in each switch.

See the [“Capturing Frames Locally”](#) section on page 29-9.

Remote Capture Daemon

This daemon is the server end of the remote capture component. The Ethereal analyzer running on a host is the client end. They communicate with each other using the Remote Capture Protocol (RPCAP). RPCAP uses two end points, a TCP-based control connection and a TCP or UDP-based data connection based on TCP (default) or UDP. The control connection is used to remotely control the captures (start or stop the capture, or specify capture filters). Remote capture can only be performed to explicitly configured hosts. This technique prevents unauthorized machine in the network from snooping on the control traffic in the network.

RPCAP supports two setup connection modes based on fire wall restrictions.

- **Passive mode (default)**—The configured host initiates connection to the switch. Multiple hosts can be configured to be in passive mode and multiple hosts can be connected and receive remote captures at the same time.
- **Active mode**—The switch initiates the connection to a configured host—one host at a time.

Using capture filters, you can limit the amount of traffic that is actually sent to the client. Capture filters are specified at the client end—on Ethereal, not on the switch.

See the “Sending Captures to Remote IP Addresses” section on page 29-11.

GUI-Based Client

The Ethereal software runs on a host, such as a PC or workstation, and communicates with the remote capture daemon. This software is available in the public domain from <http://www.ethereal.com>. Since Ethereal has a GUI front-end, it supports a rich functionality such as colorized display, graphical assists in defining filters, and searching for specific frames. These features are documented on Ethereal’s web site.

While remote capture via Ethereal supports capturing and decoding Fibre Channel frames from a Cisco MDS 9000 Family switch, the host running Ethereal does not require a Fibre Channel connection to the switch. The remote capture daemon running on the switch sends the captured frames over the out-of-band Ethernet management port. This capability allows you to capture and decode Fibre Channel frames from your desktop or laptop.

See the “Display Filters” section on page 29-12.

Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer by issuing the **fcanalyzer local** or **fcanalyzer remote** commands in configuration mode.

- **Local capture**—The command setting to enable a local capture cannot be saved to persistent storage or synchronized to standby.
- **Remote capture**—The command setting to enable a remote capture can be saved to persistent storage using the **copy** command. It can be synchronized to the standby supervisor module and a stateless restart can be issued, if required.

To use the Cisco Fabric Analyzer feature, traffic should be flowing to or from the supervisor module.

Capturing Frames Locally

Launches the textual version on the analyzer directly on the console screen. The capture can also be saved on the local file system.

To capture frames locally, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
	Note The options within Step 2 may be performed in any order.	

	Command	Purpose
Step 2	switch(config)# fcanalyzer local Capturing on eth2 switch(config)#	Begins capturing the frames locally (supervisor module).
	switch(config)# fcanalyzer local brief Capturing on eth2 switch(config)#	Displays the protocol summary in a brief format.
	switch(config)# fcanalyzer local display-filter SampleF Capturing on eth2	Displays the filtered frames.
	switch(config)# fcanalyzer local limit-frame-size 64 Capturing on eth2 switch(config)#	Limits the size of the frame capture to the first 64 bytes. The allowed range is 64 to 65536 bytes.
	switch(config)# fcanalyzer local limit-captured-frames 10 Capturing on eth2 switch(config)#	Limits the number of frames captured to 10. The allowed range is 0 to 2147483647 frames and the default is 100 frames. Use 0 if you do not want to limit the number of captured frames.
	Note Press Ctrl-c to stop a capture. Otherwise, the capture stops automatically after capturing 100 frames. You can change this default using the fcanalyzer local limit-captured-frames <i>number</i> command.	
Step 3	switch(config)# fcanalyzer local write volatile:sample Capturing on eth2 switch(config)#	Saves the captured frames to a specified file (sample) in the volatile: directory. Note Optionally, you can save the specified file to the slot0: directory.
	Note The final filename that is the capture file will be called either SampleFile_00000_<dateandtime> or SampleFile_00001_<dateandtime>. For example, “SampleFile_00000_20021110223833” or “SampleFile_00001_20021110243833”. The maximum size of a file that can be written to is 10MB.	

Sending Captures to Remote IP Addresses



Caution

You must use the eth2 interface to capture control traffic on a supervisor-module.

To capture frames remotely, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcanalyzer remote 10.21.0.3 switch(config)#	Configures the remote IP address (10.21.0.3) to which the captured frames will be sent.
	switch(config)# fcanalyzer remote 10.21.0.3 active switch(config)#	Enables active mode (passive is the default) with the remote host. Ethereal is assumed to be running when the capture is performed. The switch tries to connect forever unless a capture stop instruction is sent from Ethereal.
	switch(config)# fcanalyzer remote 10.21.0.3 active 1 switch(config)#	Enables the active mode for a specified port. The valid port range is 1 to 65535.

To capture remote traffic, use one of the following options:

- To specify the capture interface in Ethereal as the remote device:

```
rpcap://<ipaddress or switch hostname>/eth2
```

For example:

```
rpcap://cp-16/eth2
rpcap://17.2.1.1/eth2
```

- The capture interface can be specified either in the capture dialog box or using the -i option at the command line when invoking Ethereal.

```
ethereal -i rpcap://<ipaddress|hostname>[:<port>]/<interface>
```

For example:

```
ethereal -i rpcap://172.22.1.1/eth2
```

or

```
ethereal -i rpcap://customer-switch.customer.com/eth2
```



Note

For example, in a Windows 2000 setup, click **Start** on your desktop and select **Run...** In the resulting Run window, type the required command line option in the Open field.

Clearing Configured fcanalyzer Information

Use the **clear fcanalyzer** command to clear the entire list of configured hosts. Note that the existing connections are not terminated.

Viewing Display Filters Information

Use the **show fcanalyzer** command to display the list of hosts configured for a remote capture. See [Example 29-3](#).

Example 29-3 Displays Configured Hosts

```
switch# show fcanalyzer
PassiveClient = 10.21.0.3
PassiveClient = 10.21.0.3
ActiveClient = 10.21.0.3, DEFAULT
```

**Note**

The DEFAULT in the ActiveClient line indicates that the default port is used.

Display Filters

You can selectively view captured frames by using the display filters feature. For example, instead of viewing all the frames from a capture, you may only want to view ELP request frames. This feature only limits the captured view—it does not affect the captured or the saved frames. Procedures to specify, use, and save display filters are already document in the Ethereal web site (<http://www.ethereal.com>). Some examples of how you can use this feature are provided below:

- To view all packets in a specified VSAN, use this expression:

```
mdshdr.vsan == 2
```

- To view all SW_ILS frames, use this expression:

```
fcswils
```

- To view class F frames, use this expression:

```
mdshdr.sof == SOFf
```

- To view all FSPF frames, use this expression:

```
swils.opcode == JLO || swils.opcode == LSU || swils.opcode == LSA
```

- To view all FLOGI frames, use this expression:

```
fcels.opcode == FLOGI
```

- To view all FLOGI frames in VSAN 1, use this expression:

```
fcels.opcode == FLOGI && mdshdr.vsan == 2
```

- To view all name server frames, use this expression:

```
dns
```

Defining Display Filters

Display filters limit the frames that can be displayed, but not what is captured (similar to any view command). The filters to be displayed can be defined in multiple ways in the GUI application:

- Auto-definition
- Manual definition
- Assisted manual definition
- Only manual definition in local capture
- No assists

Regardless of the definition, each filter must be saved and identified with a name.



Note

This GUI-assisted feature is part of Ethereal and you can obtain more information from <http://www.ethereal.com>.

Displaying Filters Examples

Some examples of using display filters with the Fabric Analyzer local are provided in this section. The **brief** option is used in all examples to restrict the size of the output. See [Example 29-4](#).

Example 29-4 Displays Only Fabric Login Server Traffic on VSAN 1

```
switch(config)# fcanalyzer local brief display-filter
mdshdr.vsan==0x01)&&((fc.d_id==ff.ff.fe)or(fc.s_id==ff.ff.fe))
Capturing on eth2
8.904145 00.00.00 -> ff.ff.fe FC ELS 1 0x28f8 0xffff 0x3 -> 0xf FLOGI
8.918164 ff.ff.fe -> 79.03.00 FC ELS 1 0x28f8 0x12c6 0xff -> 0x0 ACC (FLOGI)
```

You can trace all frames to and from a particular N port device. For example, to observe RSCNs from Fabric Controller and registration and/or query requests to the Name Server. See [Example 29-5](#).



Note

The filter requires prior knowledge of the FC ID that is assigned to the N port. Issue the **show flogi database interface** command before running fcanalyzer to obtain the FC ID. In this example, the N port FC ID is 79.03.00.

Example 29-5 Displays All Traffic for a Particular N Port on VSAN 1

```
switch(config)# fcanalyzer local brief
display-filter(mdshdr.vsan==0x01)&&((fc.d_id==79.03.00)or(fc.s_id==79.03.00))
Capturing on eth2
8.699162 ff.ff.fe -> 79.03.00 FC ELS 1 0x35b8 0x148e 0xff -> 0x0 ACC (FLOGI)
8.699397 79.03.00 -> ff.ff.fc FC ELS 1 0x35d0 0xffff 0x3 -> 0xf PLOGI
8.699538 ff.ff.fc -> 79.03.00 FC ELS 1 0x35d0 0x148f 0xff -> 0x0 ACC (PLOGI)
8.699406 79.03.00 -> ff.ff.fd FC ELS 1 0x35e8 0xffff 0x3 -> 0xf SCR
8.700179 79.03.00 -> ff.ff.fc dNS 1 0x3600 0xffff 0x3 -> 0xf GNN_FT
8.702446 ff.ff.fd -> 79.03.00 FC ELS 1 0x35e8 0x1490 0xff -> 0x0 ACC (SCR)
8.704210 ff.ff.fc -> 79.03.00 dNS 1 0x3600 0x1491 0xff -> 0x0 ACC (GNN_FT)
8.704383 79.03.00 -> ff.ff.fc dNS 1 0x3618 0xffff 0x3 -> 0xf GPN_ID
8.707857 ff.ff.fc -> 79.03.00 dNS 1 0x3618 0x1496 0xff -> 0x0 ACC (GPN_ID)
```

The VSAN ID is specified in hex. See [Example 29-6](#).

Example 29-6 Displays All Traffic for a Specified VSAN

```
switch(config)# fcanalyzer local brief display-filter mdshdr.vsan==0x03e7
Capturing on eth2
12.762577 ff.ff.fd -> ff.ff.fd SW_ILS 999 0xb2c 0xffff 0x1 -> 0xf HLO
12.762639 ff.ff.fd -> ff.ff.fd FC 999 0xb2c 0xd32 0xff -> 0x0 Link Ctl, ACK1
13.509979 ff.ff.fd -> ff.ff.fd SW_ILS 999 0xd33 0xffff 0xff -> 0x0 HLO
13.510918 ff.ff.fd -> ff.ff.fd FC 999 0xd33 0xb2d 0x1 -> 0xf Link Ctl, ACK1
14.502391 ff.fc.64 -> ff.fc.70 SW_ILS 999 0xd34 0xffff 0xff -> 0x0 SW_RSCN
14.502545 ff.ff.fd -> 64.01.01 FC ELS 999 0xd35 0xffff 0xff -> 0x0 RSCN
14.502804 64.01.01 -> ff.ff.fd FC ELS 999 0xd35 0x215 0x0 -> 0xf ACC (RSCN)
14.503387 ff.fc.70 -> ff.fc.64 FC 999 0xd34 0xb2e 0x1 -> 0xf Link Ctl, ACK1
14.503976 ff.fc.70 -> ff.fc.64 SW_ILS 999 0xd34 0xb2e 0x1 -> 0xf SW_ACC (SW_RSCN)
14.504025 ff.fc.64 -> ff.fc.70 FC 999 0xd34 0xb2e 0xff -> 0x0 Link Ctl, ACK1
```

By excluding FSPF hellos and ACK1, you can focus on the frames of interest. See [Example 29-7](#).

Example 29-7 Displays All VSAN 1 Traffic Excluding FSPF hellos and ACK1 Frames.

```
switch(config)# fcan lo bri dis
(mdshdr.vsan==0x01)&&not((swils.opcode==0x14)or(fc.r_ctl==0xc0))
Capturing on eth2
10.589934 ff.fc.79 -> ff.fc.7a FC-FCS 1 0x1b23 0xffff 0xff -> 0x0 GCAP
10.591253 ff.fc.7a -> ff.fc.79 FC-FCS 1 0x1b23 0x2f70 0x4 -> 0xf MSG_RJT (GCAP)
25.277981 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1b27 0xffff 0xff -> 0x0 SW_RSCN
25.278050 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1b28 0xffff 0xff -> 0x0 SW_RSCN
25.279232 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1b28 0xadd7 0x5 -> 0xf SW_ACC (SW_RSCN)
25.280023 ff.fc.7a -> ff.fc.79 Unzoned NS 1 0x3b2b 0xffff 0x5 -> 0xf GE_PT
25.280029 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1b27 0x2f71 0x4 -> 0xf SW_ACC (SW_RSCN)
25.282439 ff.fc.79 -> ff.fc.7a dNS 1 0x3b2b 0x1b29 0xff -> 0x0 RJT (GE_PT)
38.249966 00.00.00 -> ff.ff.fe FC ELS 1 0x36f0 0xffff 0x3 -> 0xf FLOGI
38.262622 ff.ff.fe -> 79.03.00 FC ELS 1 0x36f0 0x1b2b 0xff -> 0x0 ACC (FLOGI)
38.262844 79.03.00 -> ff.ff.fc FC ELS 1 0x3708 0xffff 0x3 -> 0xf PLOGI
38.262984 ff.ff.fc -> 79.03.00 FC ELS 1 0x3708 0x1b2c 0xff -> 0x0 ACC (PLOGI)
38.262851 79.03.00 -> ff.ff.fd FC ELS 1 0x3720 0xffff 0x3 -> 0xf SCR
38.263514 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1b2e 0xffff 0xff -> 0x0 SW_RSCN
38.263570 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1b2f 0xffff 0xff -> 0x0 SW_RSCN
38.263630 79.03.00 -> ff.ff.fc dNS 1 0x3738 0xffff 0x3 -> 0xf GNN_FT
38.263884 ff.ff.fd -> 79.03.00 FC ELS 1 0x3720 0x1b2d 0xff -> 0x0 ACC (SCR)
38.264066 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1b2f 0xaddf 0x5 -> 0xf SW_ACC (SW_RSCN)
38.264417 ff.fc.89 -> ff.fc.79 dNS 1 0xade0 0xffff 0x5 -> 0xf GE_ID
38.264585 ff.fc.79 -> ff.fc.89 dNS 1 0xade0 0x1b31 0xff -> 0x0 ACC (GE_ID)
38.265132 ff.ff.fc -> 79.03.00 dNS 1 0x3738 0x1b30 0xff -> 0x0 ACC (GNN_FT)
38.265210 ff.fc.7a -> ff.fc.79 Unzoned NS 1 0x3b2f 0xffff 0x5 -> 0xf GE_PT
38.265414 79.03.00 -> ff.ff.fc dNS 1 0x3750 0xffff 0x3 -> 0xf GPN_ID
38.265502 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1b2e 0x2f73 0x4 -> 0xf SW_ACC (SW_RSCN)
38.267196 ff.fc.79 -> ff.fc.7a dNS 1 0x3b2f 0x1b32 0xff -> 0x0 ACC (GE_PT)
```

Use this command to focus on TE port initialization. This example allows two VSANs on the TE port and the port VSAN is 666. Hence the ELP, ESC and EPP (0x71) go out on VSAN 666. Once the EPP negotiation is complete we see EFP, DIA, RDI, MR, FSPF, and other updates flow for each allowed VSAN. See [Example 29-8](#).

Example 29-8 Displays SW_ILS Traffic between Fabric Controllers for all VSANs and Exclude FSPF hellos and ACK1 frames.

```
switch(config)# fcan lo bri dis
((fc.s_id==ff.ff.fd)&&(fc.type==0x22))&&not((swils.opcode==0x14))
Capturing on eth2
```

```

20.573225 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200c 0xffff 0xe -> 0xf ELP
20.574021 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200c 0xacc4 0xff -> 0x0 SW_ACC (ELP)
20.606020 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200d 0xffff 0xe -> 0xf ESC
20.606232 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200d 0xacc5 0xff -> 0x0 SW_ACC (ESC)
20.606665 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200e 0xffff 0xe -> 0xf 0x71
20.608768 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x200e 0xacc6 0xff -> 0x0 SW_ACC (0x71)
20.615346 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc7 0xffff 0xff -> 0x0 0x71
20.620330 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc7 0x200f 0xe -> 0xf SW_ACC (0x71)
20.623028 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2010 0xffff 0xe -> 0xf EFP
20.624681 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc9 0xffff 0xff -> 0x0 EFP
20.624974 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2010 0xacc8 0xff -> 0x0 SW_ACC (EFP)
20.625133 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1939 0xffff 0xff -> 0x0 EFP
20.626393 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacc9 0x2011 0xe -> 0xf SW_ACC (EFP)
20.627185 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0b 0xffff 0xe -> 0xf EFP
20.627479 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1939 0xab0a 0xe -> 0xf SW_ACC (EFP)
20.627773 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0c 0xffff 0xe -> 0xf DIA
20.631106 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0b 0x193a 0xff -> 0x0 SW_ACC (EFP)
20.631432 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0d 0xffff 0xe -> 0xf MR
20.631567 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193c 0xffff 0xff -> 0x0 DIA
20.631974 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0c 0x193b 0xff -> 0x0 SW_ACC (DIA)
20.631938 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193c 0xab0e 0xe -> 0xf SW_ACC (DIA)
20.639262 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193e 0xffff 0xff -> 0x0 MR
20.640417 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x193e 0xab0f 0xe -> 0xf SW_ACC (MR)
20.640598 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab0d 0x193d 0xff -> 0x0 SW_ACC (MR)
20.646950 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab14 0xffff 0xe -> 0xf LSU
20.647256 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1944 0xffff 0xff -> 0x0 LSU
20.647996 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1945 0xffff 0xff -> 0x0 LSU
20.648367 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1946 0xffff 0xff -> 0x0 LSA
20.648476 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab17 0xffff 0xe -> 0xf LSU
20.648916 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab19 0xffff 0xe -> 0xf LSA
20.649210 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab1a 0xffff 0xe -> 0xf LSA
20.659781 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x194a 0xffff 0xff -> 0x0 LSA
20.660535 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab1d 0xffff 0xe -> 0xf LSU
20.660649 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x194c 0xffff 0xff -> 0x0 LSU
20.660683 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab1e 0xffff 0x5 -> 0xf LSU
20.661006 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x194e 0xffff 0xff -> 0x0 LSU
20.664994 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab22 0xffff 0xe -> 0xf LSA
20.665341 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab24 0xffff 0x5 -> 0xf LSU
20.665645 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab25 0xffff 0x5 -> 0xf LSA
20.666115 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1952 0xffff 0xff -> 0x0 LSA
20.666445 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1953 0xffff 0xff -> 0x0 LSU
20.666994 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1954 0xffff 0xff -> 0x0 LSA
20.667423 ff.ff.fd -> ff.ff.fd SW_ILS 1 0xab2a 0xffff 0x5 -> 0xf LSA
20.667715 ff.ff.fd -> ff.ff.fd SW_ILS 1 0x1956 0xffff 0xff -> 0x0 LSA
30.525363 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2012 0xffff 0xe -> 0xf DIA
30.525596 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2012 0xacca 0xff -> 0x0 SW_ACC (DIA)
30.525959 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xaccb 0xffff 0xff -> 0x0 RDI
30.526736 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xaccb 0x2013 0xe -> 0xf SW_ACC (RDI)
30.527032 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2014 0xffff 0xe -> 0xf EFP
30.527662 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2014 0xacc 0xff -> 0x0 SW_ACC (EFP)
30.533157 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2015 0xffff 0xe -> 0xf MR
30.534159 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacce 0xffff 0xff -> 0x0 MR
30.534440 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2015 0xaccd 0xff -> 0x0 SW_ACC (MR)
30.534791 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacce 0x2016 0xe -> 0xf SW_ACC (MR)
30.540883 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x201b 0xffff 0xe -> 0xf LSU
30.541068 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacd4 0xffff 0xff -> 0x0 LSU
30.541704 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacd5 0xffff 0xff -> 0x0 LSA
30.541981 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacd6 0xffff 0xff -> 0x0 LSU
30.542087 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x201e 0xffff 0xe -> 0xf LSA
30.542381 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2020 0xffff 0xe -> 0xf LSU
30.542675 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2021 0xffff 0xe -> 0xf LSU
30.542969 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2022 0xffff 0xe -> 0xf LSA
30.543226 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacdb 0xffff 0xff -> 0x0 LSU
30.543614 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2024 0xffff 0xe -> 0xf LSA

```

```

30.543751 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacdd 0xffff 0xff -> 0x0 LSA
30.544004 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacde 0xffff 0xff -> 0x0 LSA
30.544522 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xacdf 0xffff 0xff -> 0x0 LSU
30.544553 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x2027 0xffff 0xe -> 0xf LSU
30.550961 ff.ff.fd -> ff.ff.fd SW_ILS 666 0xace7 0xffff 0xff -> 0x0 LSA
30.550988 ff.ff.fd -> ff.ff.fd SW_ILS 666 0x202f 0xffff 0xe -> 0xf LSA

```

This example focuses on zone server changes. Prior knowledge of the domain controller ID is required. The switch domain ID where the fcanalyzer is run is x79, the domain controller is FF.FC.79. See [Example 29-9](#).

Example 29-9 Display SW_ILS traffic between Fabric Domain Controllers for VSAN 1

```

switch(config)# fcan lo bri dis
mdshdr.vsan==0x01&&(fc.type==0x22)&&((fc.d_id==ff.fc.79)or(fc.s_id==ff.fc.79))
Capturing on eth2
64.053927 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e15 0xffff 0xff -> 0x0 ACA
64.053995 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e16 0xffff 0xff -> 0x0 ACA
64.054599 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e16 0xb1e2 0x5 -> 0xf SW_ACC (ACA)
64.054747 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e15 0x3037 0x4 -> 0xf SW_ACC (ACA)
64.057643 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e17 0xffff 0xff -> 0x0 SFC
64.057696 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e18 0xffff 0xff -> 0x0 SFC
64.058788 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e17 0x3038 0x5 -> 0xf SW_ACC (SFC)
64.059288 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e18 0xb1e3 0x5 -> 0xf SW_ACC (SFC)
64.062011 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e19 0xffff 0xff -> 0x0 UFC
64.062060 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e1a 0xffff 0xff -> 0x0 UFC
64.073513 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e19 0x3039 0x5 -> 0xf SW_ACC (UFC)
64.765306 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e1a 0xb1e4 0x5 -> 0xf SW_ACC (UFC)
64.765572 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1e1b 0xffff 0xff -> 0x0 RCA
64.765626 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1e1c 0xffff 0xff -> 0x0 RCA
64.766386 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1e1b 0x303a 0x4 -> 0xf SW_ACC (RCA)
64.766392 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1e1c 0xb1e5 0x5 -> 0xf SW_ACC (RCA)

```

Capture Filters

You can limit what frames are captured by using the capture filters feature in a remote capture. This feature limits the frames that are captured and sent from the remote switch to the host. For example, you can capture only class F frames. Capture filters is useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restricts a capture to the specified frames. No other frames will be visible until you specify a completely new capture.

The syntax for capture filter is different from the syntax for display filters. Capture filters use the Berkeley Packet Filter (BPF) library that is used in conjunction with the libpcap freeware. The list of all valid Fibre Channel capture filter fields are provided later in this section.

Procedures to configure capture filters are already documented in the Ethereal web site (<http://www.ethereal.com>). Some examples of how you can use this feature are provided below:

- To capture frames only on a specified VSAN, use this expression:

```
vsan = 1
```

- To capture only class F frames, use this expression:

```
class_f
```

- To capture only class Fibre Channel ELS frames, use this expression:

```
els
```


- To capture only name server frames, use this expression:
dns
- To capture only SCSI command frames, use this expression:
fcpcmd

**Note**

This feature is part of libpcap and you can obtain more information from <http://www.tcpdump.org>.

Permitted Capture Filters

```

o vsan
o src_port_idx
o dst_port_idx
o sof
o r_ctl
o d_id
o s_id
o type
o seq_id
o seq_cnt
o ox_id
o rx_id
o els
o swils
o fcp_cmd    (FCP Command frames only)
o fcp_data   (FCP data frames only)
o fcp_rsp    (FCP response frames only)
o class_f
o bad_fc
o els_cmd
o swils_cmd
o fcp_lun
o fcp_task_mgmt
o fcp_scsi_cmd
o fcp_status
o gs_type    (Generic Services type)
o gs_subtype (Generic Services subtype)
o gs_cmd
o gs_reason
o gs_reason_expl
o dns        (name server)
o udns       (unzoned name server)
o fcs        (fabric configuration server)
o zs         (zone server)
o fc         (use as fc[x:y] where x is offset and y is length to compare)
o els        (use as els[x:y] similar to fc)
o swils      (use as swils[x:y] similar to fc)
o fcp        (use as fcp[x:y] similar to fc)
o fcct       (use as fcct[x:y] similar to fc)

```

Configuring World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch’s supervisor module, assigns WWNs to each switch. This WWN is independent of other WWNs on each switch. This centralized control of WWN has the following advantages:

- Efficient sharing of WWN space
- Centralized support across switches

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see [Table 29-1](#)).

Table 29-1 Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits



Caution

Changes to the worldwide names should be made by an administrator or individual who is completely familiar with switch operations.

Configuring a Secondary MAC Address

To allocate secondary MAC addresses, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# wwn secondary-mac 00:99:55:77:55:55 range 64 This command CANNOT be undone. Please enter the BASE MAC ADDRESS again: 00:99:55:77:55:55 Please enter the mac address RANGE again: 64 From now on WWN allocation would be based on new MACs. Are you sure? (yes/no) no You entered: no. Secondary MAC NOT programmed switch(config)#	Configures the secondary MAC address. This command cannot be undone.

Displaying WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. See Examples 29-10 to 29-12.

Example 29-10 Displays the Status of All WWNs

```
switch# show wwn status
      Type 1 WWNs: Configured:      64 Available:      48 (75%) Resvd.: 16
      Types 2 & 5 WWNs: Configured: 524288 Available: 450560 (85%) Resvd.: 73728
NKAU & NKCR WWN Blks: Configured:   1760 Available:   1760 (100%)
      Alarm Status:      Type1:    NONE Types 2&5:    NONE
```

Example 29-11 Displays Specified Block ID Information:

```
switch# show wwn status block-id 51
WWNs in this block: 21:00:ac:16:5e:52:00:03 to 21:ff:ac:16:5e:52:00:03
Num. of WWNs:: Configured: 256 Allocated:    0 Available: 256
Block Allocation Status: FREE
```

Example 29-12 Displays the WWN for a Specific Switch

```
switch# show wwn switch
Switch WWN is 20:00:ac:16:5e:52:00:00
```

Allocating Flat FC IDs

Fibre Channel standards require a unique FC ID to be allocated to a N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Based on the assigned FC ID, some HBAs assume that no other ports have the same area bits and domain. When a target is assigned with a FC ID that has the same area bits, but different port bits, the HBA fails to discover these targets. To isolate these HBAs in a separate area, switches in the Cisco MDS 9000 Family follow a different FC ID allocation scheme. By default, the FC ID allocation mode is **auto**. In the **auto** mode, only HBAs without interop issues are assigned FCIDs with specific ports bits. All other HBAs are assigned FC IDs with a whole area (port bits set to 0).

The three options to allocate FCID are auto (default), none, and flat.

To allocate flat FC IDs, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcinterop fcid-allocation none	Allocates one area to the N port attached to an F port.
	switch(config)# fcinterop fcid-allocation flat	Allocates a single FC ID to the N port. This option is generally used to conserve FC ID usage.
	switch(config)# fcinterop fcid-allocation auto	Intelligently assigns flat FC ID to N ports which can interoperate in flat mode, otherwise assigns full area to all other ports. This is the default.



Caution

Changes to FC IDs should be made by an administrator or individual who is completely familiar with switch operations.

Enabling Loop Monitoring

When a disk is removed from a loop port, the loop stays active based on the bypass circuit. Thus the disk removal is not known until you try to communicate with the disk. To detect such removals, the disks can be polled periodically (every 20 seconds) using the **fcinterop loop-monitor** command. This command enables loop polling for FL ports in a Cisco MDS 9000 Family switch. By default, the **fcinterop loop-monitor** command is disabled.

To enable the loop monitoring feature, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcinterop loop-monitor	Enables the loop monitoring feature.
	switch(config)# no fcinterop loop-monitor	Disables (default) the loop monitoring feature and reverts the switch to the factory defaults.



Caution

Changes to the loop monitoring feature should be made by an administrator or individual who is completely familiar with switch operations.

Configuring the Switch for Interoperability

Interoperability enables multiple vendors' products come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way thus resulting in interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provide the product with a more aimable standards compliant implementation.

Table 29-2 lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

Table 29-2 Changes in Switch Behavior When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be setup statically (the MDS will only accept one domain ID, if it doesn't get that domain ID it isolates itself from the fabric), or preferred. (If it doesn't get its requested domain ID, it accepts any assigned domain ID.)
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are:
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.
Default zone	The default zone behavior of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone), may change.
Zoning attributes	Zones may be limited to the pWWN and other proprietary zoning methods (physical port number), may be eliminated. Note Brocade uses the cfgsave command to save fabric-wide zoning configuration. This command doesn't have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family using the copy running start command.
Zone propagation	Some vendors do not pass the full zone configuration (zoneset) to other switches, only the active zoneset gets passed. Verify that the active zoneset or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN.

Table 29-2 Changes in Switch Behavior When Interoperability Is Enabled (continued)

Switch Feature	Changes if Interoperability Is Enabled
TE ports and PortChannels	TE ports and Port-Channels cannot be used to connect MDS to non-MDS switches. Only E ports can be used to connect to non-MDS switches. TE ports and PortChannels can still be used to connect an MDS to other MDS switches even when in interop mode.
FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to loadbalance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing Domain IDs.
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.
IVR	IVR-enabled VSANs can only be configured in either no interop mode or interop 1 mode.

Configuring Interoperability

The **interop** mode in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.



Note

Brocade's `msplmgmtdeactivate` command must explicitly be run prior to connect from a Brocade switch to either Cisco MDS 9000 Family switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco MDS 9000 Family switches or McData switches do not understand. Rejecting these frames, causes the common E ports to become isolated.

To configure interoperability in any switch in the Cisco MDS 9000 Family, follow these steps:

- Step 1** Place the VSAN of the E ports (s) that connect to the OEM switch in interoperability mode.

```
switch# config t
switch(config)# vsan database
switch (config-vsan-db)# vsan 1 interop 1
```

- Step 2** Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).



Note

This is an limitation imposed by the McData switches.

```
switch# config t
switch(config)# fcdomain domain 100 preferred vsan 1
```

In Cisco MDS 9000 switches, the default is to request an ID from the principal switch. If the **preferred** option is used, Cisco MDS 9000 switches request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the **static** option is used, the Cisco MDS 9000 switches does not join the fabric unless the principal switch agrees, and assigns the requested ID.



Note When changing the Domain ID, the FC IDs assigned to N ports will also change.

Step 3 Change the Fibre Channel timers (if they have been changed from the system defaults).



Note The MDS 9000, Brocade, and McData FC Error Detect (ED_TOV) and Resource Allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

```
switch# config t
switch(config)# fctimer e_d_tov ?
    <1000-100000>  E_D_TOV in milliseconds (1000-100000)
switch(config)# fctimer r_a_tov ?
    <5000-100000>  R_A_TOV in milliseconds (5000-100000)
```

Step 4 When making changes to the domain, you may or may not need to restart the MDS domain manager function for the altered VSAN.

- Force a fabric reconfiguration with the **disruptive** option.

```
switch(config)# fcdomain restart disruptive vsan 1
```

or

- Don't force a fabric reconfiguration

```
switch(config)# fcdomain restart vsan 1
```

Verifying Interoperating Status

This section highlights the commands used to verify if the fabric is up and running in interoperability mode.

To verify the resulting status of issuing the interoperability command in any switch in the Cisco MDS 9000 Family, follow these steps:

Step 1 Use the **show version** command to verify the version.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Andiamo Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.
```

Software

```

BIOS:          version 1.0.8
loader:        version 1.1(2)
kickstart:     version 2.0(1) [build 2.0(0.6)] [gdb]
system:        version 2.0(1) [build 2.0(0.6)] [gdb]

BIOS compile time:      08/07/03
kickstart image file is: bootflash:///m9500-sf1ek9-kickstart-mzg.2.0.0.6.bin
kickstart compile time: 10/25/2010 12:00:00
system image file is:   bootflash:///m9500-sf1ek9-mzg.2.0.0.6.bin
system compile time:    10/25/2020 12:00:00

```

Hardware

```

RAM 1024584 kB

bootflash: 1000944 blocks (block size 512b)
slot0:      0 blocks (block size 512b)

172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
Reason: Reset Requested by CLI command reload
System version: 2.0(0.6)
Service:

```

Step 2 Use the **show interface brief** command to verify if the interface states are as required by your configuration.

```

switch# show int brief

```

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	Oper Mode	Oper Speed (Gbps)	Port-channel
fc2/1	1	auto	on	up	E	2	--
fc2/2	1	auto	on	up	E	2	--
fc2/3	1	auto	on	fcotAbsent	--	--	--
fc2/4	1	auto	on	down	--	--	--
fc2/5	1	auto	on	down	--	--	--
fc2/6	1	auto	on	down	--	--	--
fc2/7	1	auto	on	up	E	1	--
fc2/8	1	auto	on	fcotAbsent	--	--	--
fc2/9	1	auto	on	down	--	--	--
fc2/10	1	auto	on	down	--	--	--

Step 3 Use the **show run** command to verify if you are running the desired configuration.

```

switch# show run
Building Configuration...

interface fc2/1
no shutdown

interface fc2/2
no shutdown

interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
interface fc2/7
no shutdown

interface fc2/8
interface fc2/9

```



```

interface fc2/10

<snip>

interface fc2/32

interface mgmt0
ip address 6.1.1.96 255.255.255.0
switchport encap default
no shutdown

vsan database
vsan 1 interop

boot system bootflash:/m9500-system-253e.bin sup-1
boot kickstart bootflash:/m9500-kickstart-253e.bin sup-1
boot system bootflash:/m9500-system-253e.bin sup-2
boot kickstart bootflash:/m9500-kickstart-253e.bin sup-2
callhome

fcdomain domain 100 preferred vsan 1

ip route 6.1.1.0 255.255.255.0 6.1.1.1
ip routing
line console
  databits 5
  speed 110
logging linecard
ssh key rsa 512 force
ssh server enable
switchname MDS9509
username admin password 5 $1$Li8/fBYX$SNc72.xt4nTXpSnR9OUFB/ role network-admin

```

Step 4 Use the **show vsan** command to verify if the interoperability mode is active.

```

switch# show vsan 1
vsan 1 information
  name:VSAN0001 stalactites
  interoperability mode:yes <----- verify mode
  loadbalancing:src-id/dst-id/oxid
  operational state:up

```

Step 5 Use the **show fcdomain vsan** command to verify the domain ID.

```

switch# show fcdomain vsan 1
The local switch is a Subordinated Switch.

Local switch run time information:
  State: Stable
  Local switch WWN: 20:01:00:05:30:00:51:1f
  Running fabric name: 10:00:00:60:69:22:32:91
  Running priority: 128
  Current domain ID: 0x64(100) <-----verify domain id

Local switch configuration information:
  State: Enabled
  Auto-reconfiguration: Disabled
  Contiguous-allocation: Disabled
  Configured fabric name: 41:6e:64:69:61:6d:6f:21
  Configured priority: 128
  Configured domain ID: 0x64(100) (preferred)

Principal switch run time information:
  Running priority: 2

```

Interface	Role	RCF-reject
fc2/1	Downstream	Disabled
fc2/2	Downstream	Disabled
fc2/7	Upstream	Disabled

Step 6 Use the **show fcdomain domain-list vsan** command to verify the local principal switch status.

```
switch# show fcdomain domain-list vsan 1

Number of domains: 5
Domain ID          WWN
-----
0x61(97)           10:00:00:60:69:50:0c:fe
0x62(98)           20:01:00:05:30:00:47:9f
0x63(99)           10:00:00:60:69:c0:0c:1d
0x64(100)          20:01:00:05:30:00:51:1f [Local]
0x65(101)          10:00:00:60:69:22:32:91 [Principal]
-----
```

Step 7 Use the **show fspf internal route vsan** command to verify the next hop and destination for the switch.

```
switch# show fspf internal route vsan 1

FSPF Unicast Routes
-----
VSAN Number  Dest Domain  Route Cost  Next hops
-----
          1      0x61(97)         500      fc2/2
          1      0x62(98)        1000      fc2/1
                                   fc2/2
          1      0x63(99)         500      fc2/1
          1      0x65(101)        1000      fc2/7
```

Step 8 Use the **show fcns data vsan** command to verify the name server information.

```
switch# show fcns data vsan 1
VSAN 1:
-----
FCID          TYPE  PWWN                                (VENDOR) FC4-TYPE:FEATURE
-----
0x610400      N     10:00:00:00:c9:24:3d:90 (Emulex)    scsi-fcp
0x6105dc      NL    21:00:00:20:37:28:31:6d (Seagate)   scsi-fcp
0x6105e0      NL    21:00:00:20:37:28:24:7b (Seagate)   scsi-fcp
0x6105e1      NL    21:00:00:20:37:28:22:ea (Seagate)   scsi-fcp
0x6105e2      NL    21:00:00:20:37:28:2e:65 (Seagate)   scsi-fcp
0x6105e4      NL    21:00:00:20:37:28:26:0d (Seagate)   scsi-fcp
0x630400      N     10:00:00:00:c9:24:3f:75 (Emulex)    scsi-fcp
0x630500      N     50:06:01:60:88:02:90:cb (Seagate)   scsi-fcp
0x6514e2      NL    21:00:00:20:37:a7:ca:b7 (Seagate)   scsi-fcp
0x6514e4      NL    21:00:00:20:37:a7:c7:e0 (Seagate)   scsi-fcp
0x6514e8      NL    21:00:00:20:37:a7:c7:df (Seagate)   scsi-fcp
0x651500      N     10:00:00:e0:69:f0:43:9f (JNI)

Total number of entries = 12
```



Note

The MDS Name Server shows both local and remote entries, and does not timeout the entries.

Using the show tech-support Command

The **show tech-support** command is useful when collecting a large amount of information about your switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.

The **show tech-support** command displays the output of several **show** commands at once. The output from this command will vary depending on your configuration. Use the **show tech-support** command in EXEC mode to display general information about the switch when reporting a problem.

You can choose to have detailed information for each command or even specify the output for a particular interface, module or VSAN. Each command output is separated by line and the command precedes the output.

**Note**

Explicitly set the **terminal length** command to 0 (zero) to disable auto-scrolling and enable manually scrolling. Use the **show terminal** command to view the configured terminal size. After obtaining the output of this command, remember to reset your terminal length as required (see the [“Setting the Switch’s Terminal Length”](#) section on page 2-16)

**Tip**

You can save the output of this command to a file by appending `> filename` to the **show tech-support** command (see the [“Saving Command Output to a File”](#) section on page 2-22). If you save this file, verify you have sufficient space to do so—each of these files may take about 1.8 MB. However, you can zip this file using the **gzip filename** command (see the [“Compressing and Uncompressing Files”](#) section on page 2-22). Copy the zipped file to the required location using the **copy** command and unzip the file using the **gunzip** command ([“Copying Files”](#) section on page 2-21).

The default output of the **show tech-support** command includes the output of the following commands:

- **show version**
- **show environment**
- **show module**
- **show hardware**
- **show running-config**
- **show interface**
- **show accounting log**
- **show process**
- **show process log**
- **show processes log details**
- **show flash**

Each command is discussed in both the *Cisco MDS 9000 Family Configuration Guide* and the *Cisco MDS 9000 Family Command Reference*. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide* to obtain debug processes, procedures, and examples.



Configuring Fabric Configuration Servers

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

- [About FCS, page 30-2](#)
- [Configuring FCS, page 30-3](#)
- [Displaying FCS Information, page 30-4](#)

About FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object: Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object: Each physical port in an IE corresponds to a port object. This includes the switch ports (xE, Fx, and TL ports), and its attached Nx ports.
- Platform object: A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. platform objects reside at the edge switches of the fabric.

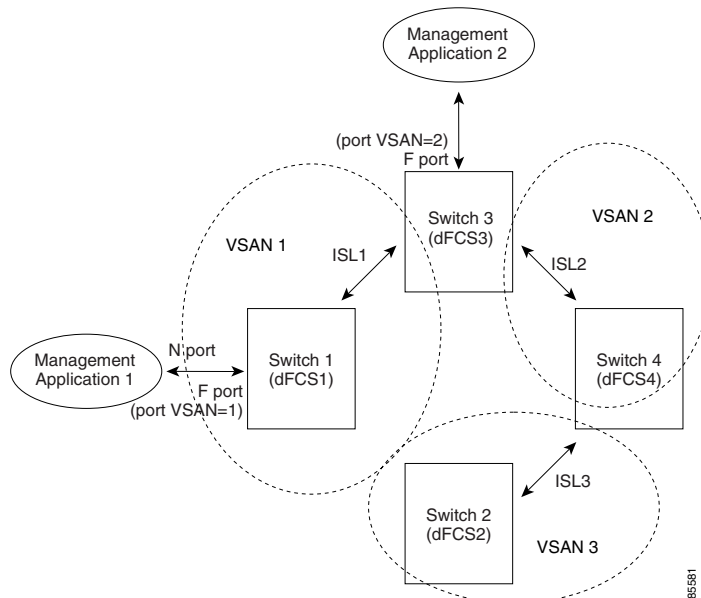
Each object has its own set of attributes and their values. A null value may also be defined for some attributes.

In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch is a part of the port VSAN in the switch port (Fx port). Hence your view of the management application is limited only to this VSAN. However information about other VSANs that this switch is part of can be obtained either through SNMP or CLI.

In Figure 2 Management Application 1 (M1) is connected through an F port with a port VSAN ID 1 and Management Application 2 (M2) is connected through an F port with a port VSAN ID 2. M1 can query FCS info of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. Note that M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 30-1 FCSs in a VSAN Environment



Significance of FCS

The significance of FCSs are as follows:

- Network Management
 - N port management application can query and obtain information about the fabric elements.
 - SNMP Manager—A SNMP Manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs supports TE and TL ports in addition to the standard F and E ports.
- FCS can maintain a group of modes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and updates it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information, and rebuild its database.
- The SNMP manager can query FCSs for all the IEs, ports, and platforms in the fabric.

Configuring FCS

Use the **fcs plat-check-global** command to specify if the platform name or node name uniqueness verification is for the entire fabric (globally) or only for locally (default) registered platforms.



Note

Set this command globally only if all switches in the fabric belong to the Cisco MDS 9000 Family.

To enable global checking of the platform name, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcs plat-check-global vsan 1 switch(config)#	Enables global checking of platform name.
	switch(config)# no fcs plat-check-global vsan 1 switch(config)#	Disables (default) global checking of platform name.

To register platform attributes, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# fcs register switch(config-fcs-register)#	Enters the FCS registration submodule
Step 3	switch(config-fcs-register)# platform name SamplePlatform vsan 1 switch(config-fcs-register-attr)#	Enters the FCS registration attributes submodule.
	switch(config-fcs-register)# no platform name SamplePlatform vsan 1 switch(config-fcs-register)#	Deletes a registered platform.

	Command	Purpose
Step 4	switch(config-fcs-register-attrib)# mgmt-addr 1.1.1.1 switch(config-fcs-register-attrib)#	Configures the platform management address.
	switch(config-fcs-register)# no mgmt-addr 1.1.1.1 switch(config-fcs-register)#	Deletes all management addresses on the platform.
Step 5	switch(config-fcs-register-attrib)# nwwn 11:22:33:44:55:66:77:88 switch(config-fcs-register-attrib)#	Configures the platform node name.
	switch(config-fcs-register)# no nwwn 11:22:33:44:55:66:77:88 switch(config-fcs-register)#	Deletes the platform node name.
Step 6	switch(config-fcs-register-attrib)# type 5 switch(config-fcs-register-attrib)#	Configures the fc-gs-3 defined platform type.
	switch(config-fcs-register)# no type 5 switch(config-fcs-register)#	Deletes the configured type and reverts the switch to its factory default of unknown type.
Step 7	switch(config-fcs-register-attrib)# exit switch(config-fcs-register)#	Exits the FCS registration attributes submode
Step 8	switch(config-fcs-register)# exit switch(config)#	Exits the FCS registration submode.

Displaying FCS Information

Use the **show fcs** commands to display the status of the WWN configuration (see Example 30-1 to 30-9).

Example 30-1 Displays FCS Local Database Information

```
switch# show fcs database
FCS Local Database in VSAN: 1
-----
Switch WWN                : 20:01:00:05:30:00:16:df
Switch Domain Id          : 0x7f(127)
Switch Mgmt-Addresses     : snmp://172.22.92.58/eth-ip
                          : http://172.22.92.58/eth-ip
Fabric-Name               : 20:01:00:05:30:00:16:df
Switch Logical-Name       : 172.22.92.58
Switch Information List    : [Cisco Systems*DS-C9509*0*20:00:00:05:30:00
Switch Ports:
-----
Interface  pWWN                Type      Attached-pWWNs
-----
fc2/1      20:41:00:05:30:00:16:de    TE        20:01:00:05:30:00:20:de
fc2/2      20:42:00:05:30:00:16:de    Unknown   None
fc2/17     20:51:00:05:30:00:16:de    TE        20:0a:00:05:30:00:20:de

FCS Local Database in VSAN: 5
-----
Switch WWN                : 20:05:00:05:30:00:12:5f
Switch Domain Id          : 0xef(239)
Switch Mgmt-Addresses     : http://172.22.90.171/eth-ip
                          : snmp://172.22.90.171/eth-ip
                          : http://10.10.15.10/vsan-ip
                          : snmp://10.10.15.10/vsan-ip
Fabric-Name               : 20:05:00:05:30:00:12:5f
Switch Logical-Name       : 172.22.90.171
```



```
Switch Information List : [Cisco Systems*DS-C9509**20:00:00:05:30:00:12:5e]
Switch Ports:
```

Interface	pWWN	Type	Attached-pWWNs
fc3/1	20:81:00:05:30:00:12:5e	TE	22:01:00:05:30:00:12:9e
fc3/2	20:82:00:05:30:00:12:5e	TE	22:02:00:05:30:00:12:9e
fc3/3	20:83:00:05:30:00:12:5e	TE	22:03:00:05:30:00:12:9e

Example 30-2 Displays a List of All IEs for a Specific VSAN

```
switch# show fcs ie vsan 1
IE List for VSAN: 1
```

IE-WWN	IE-Type	Mgmt-Id
20:01:00:05:30:00:16:df	Switch (Local)	0xfffc7f
20:01:00:05:30:00:20:df	Switch (Adjacent)	0xfffc64

[Total 2 IEs in Fabric]

Example 30-3 Displays Interconnect Element Object Information for a Specific nWWN

```
switch# show fcs ie nwwn 20:01:00:05:30:00:16:df vsan 1
IE Attributes
```

```
-----
Domain-Id = 0x7f(127)
Management-Id = 0xfffc7f
Fabric-Name = 20:01:00:05:30:00:16:df
Logical-Name = 172.22.92.58
Management Address List =
    snmp://172.22.92.58/eth-ip
    http://172.22.92.58/eth-ip
Information List:
    Vendor-Name = Cisco Systems
    Model Name/Number = DS-C9509
    Release-Code = 0
```

Example 30-4 Displays Platform Information for a Specific Platform

```
switch# show fcs platform name SamplePlatform vsan 1
Platform Attributes
```

```
-----
Platform Node Names:
    11:22:33:44:55:66:77:88
Platform Type = Gateway
Platform Management Addresses:
    1.1.1.1
```

Example 30-5 Displays a List of Platforms for a Specified VSAN

```
switch# show fcs platform vsan 1
Platform List for VSAN: 1
Platform-Names
-----
SamplePlatform
[Total 1 Platforms in Fabric]
```

Example 30-6 Displays a List of Switchports in a Specified VSAN

```
switch# show fcs port vsan 24
Port List in VSAN: 24
```

```

-- IE WWN: 20:18:00:05:30:00:16:df --
-----
Port-WWN                Type      Module-Type      Tx-Type
-----
20:41:00:05:30:00:16:de  TE_Port  SFP with Serial Id  Shortwave Laser
20:51:00:05:30:00:16:de  TE_Port  SFP with Serial Id  Shortwave Laser
[Total 2 switch-ports in IE]
-- IE WWN: 20:18:00:05:30:00:20:df --
-----
Port-WWN                Type      Module-Type      Tx-Type
-----
20:01:00:05:30:00:20:de  TE_Port  SFP with Serial Id  Shortwave Laser
20:0a:00:05:30:00:20:de  TE_Port  SFP with Serial Id  Shortwave Laser
[Total 2 switch-ports in IE]

```

Example 30-7 Displays Port Information for a Specified pWWN

```

switch# show fcs port pwwn 20:51:00:05:30:00:16:de vsan 24
Port Attributes
-----
Port Type = TE_Port
Port Number = 0x1090000
Attached-Port-WWNs:
    20:0a:00:05:30:00:20:de
Port State = Online

```

Example 30-8 Displays FCS Statistics

```

switch# show fcs statistics
FCS Statistics for VSAN: 1
-----
FCS Rx Get Reqs    :2
FCS Tx Get Reqs    :7
FCS Rx Reg Reqs    :0
FCS Tx Reg Reqs    :0
FCS Rx Dereg Reqs  :0
FCS Tx Dereg Reqs  :0
FCS Rx RSCNs       :0
...
FCS Statistics for VSAN: 30
-----
FCS Rx Get Reqs    :2
FCS Tx Get Reqs    :2
FCS Rx Reg Reqs    :0
FCS Tx Reg Reqs    :0
FCS Rx Dereg Reqs  :0
FCS Tx Dereg Reqs  :0
FCS Rx RSCNs       :0
FCS Tx RSCNs       :0
...

```

Example 30-9 Displays Platform Settings for Each VSAN

```

switch# show fcs vsan
-----
VSAN    Plat Check fabric-wide
-----
0001    Yes
0010    No
0020    No
0021    No
0030    No

```



Monitoring System Processes and Logs

This chapter provides details on monitoring the health of the switch. It includes the following sections:

- [Displaying System Processes, page 31-2](#)
- [Displaying System Status, page 31-5](#)
- [Configuring Core and Log Files, page 31-6](#)
- [Configuring Kernel Core Dumps, page 31-8](#)

Displaying System Processes

Use the **show processes** command to obtain general information about all processes (see Examples 31-1 to 31-6).

Example 31-1 Displays System Processes

```
switch# show processes
PID      State  PC          Start_cnt  TTY  Process
-----  -
868      S      2ae4f33e    1          -    snmpd
869      S      2acee33e    1          -    rscn
870      S      2ac36c24    1          -    qos
871      S      2ac44c24    1          -    port-channel
872      S      2ac7a33e    1          -    ntp
-        ER      -           1          -    mdog
-        NR      -           0          -    vbuilder
```

Terms:

- PID = process ID.
- State = process state
 - D = uninterruptible sleep (usually IO)
 - R = runnable (on run queue)
 - S = sleeping
 - T = traced or stopped
 - Z = defunct (“zombie”) process
- NR = not-running
- ER = should be running but currently not-running
- PC = current program counter in hex format
- Start_cnt = how many times a process has been started (or restarted).
- TTY = terminal that controls the process. A “-” usually means a daemon not running on any particular TTY
- Process = name of the process

Example 31-2 Displays CPU Utilization Information

```
switch# show processes cpu
PID      Runtime (ms)  Invoked  uSecs  1Sec  Process
-----  -
842      3807         137001   27     0.0   sysmgr
1112     1220         67974    17     0.0   syslogd
1269     220         13568    16     0.0   fcfwd
1276     2901        15419   188     0.0   zone
1277     738         21010    35     0.0   xbar_client
1278     1159        6789    170     0.0   wwn
1279     515         67617    7      0.0   vsan
```

Terms:

- Runtime(ms) = CPU time the process has used, expressed in milliseconds
- Invoked = number of times the process has been invoked

- uSecs = microseconds of CPU time in average for each process invocation
- 1Sec = CPU utilization in percentage for the last one second

Example 31-3 Displays Process Log Information

```
switch# show processes log
Process          PID      Normal-exit  Stack-trace  Core      Log-create-time
-----
fspf             1339      N            Y            N         Jan  5 04:25
lcm              1559      N            Y            N         Jan  2 04:49
rib              1741      N            Y            N         Jan  1 06:05
```

Terms:

- Normal-exit = whether or not the process exited normally
- Stack-trace = whether or not there is a stack trace in the log
- Core = whether or not there exists a core file
- Log-create-time = when the log file got generated

Example 31-4 Displays Detail Log Information About a Process

```
switch# show processes log pid 1339
Service: fspf
Description: FSPF Routing Protocol Application

Started at Sat Jan  5 03:23:44 1980 (545631 us)
Stopped at Sat Jan  5 04:25:57 1980 (819598 us)
Uptime: 1 hours 2 minutes 2 seconds

Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 9 (no core)
CWD: /var/sysmgr/work
```

Virtual Memory:

```
CODE      08048000 - 0809A100
DATA      0809B100 - 0809B65C
BRK       0809D988 - 080CD000
STACK     7FFFFD20
TOTAL     23764 KB
```

Register Set:

```
EBX 00000005      ECX 7FFFF8CC      EDX 00000000
ESI 00000000      EDI 7FFFF6CC      EBP 7FFFF95C
EAX FFFFFFFE      XDS 8010002B      XES 0000002B
EAX 0000008E (orig) EIP 2ACE133E      XCS 00000023
EFL 00000207      ESP 7FFFF654      XSS 0000002B
```

Stack: 1740 bytes. ESP 7FFFF654, TOP 7FFFFD20

```
0x7FFFF654: 00000000 00000008 00000003 08051E95 .....
0x7FFFF664: 00000005 7FFFF8CC 00000000 00000000 .....
0x7FFFF674: 7FFFF6CC 00000001 7FFFF95C 080522CD .....\"..
0x7FFFF684: 7FFFF9A4 00000008 7FFFC34 2AC1F18C .....4.....*
```

Example 31-5 Displays All Process Log Details

```
switch# show processes log details
=====
Service: snmpd
Description: SNMP Agent

Started at Wed Jan  9 00:14:55 1980 (597263 us)
Stopped at Fri Jan 11 10:08:36 1980 (649860 us)
Uptime: 2 days 9 hours 53 minutes 53 seconds

Start type: SRV_OPTION_RESTART_STATEFUL (24)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 6 (core dumped)
CWD: /var/sysmgr/work

Virtual Memory:

      CODE      08048000 - 0804C4A0
      DATA      0804D4A0 - 0804D770
      BRK        0804DFC4 - 0818F000
      STACK      7FFFFCE0
      TOTAL      26656 KB
...
```

Example 31-6 Displays Memory Information About Processes

```
switch# show processes memory
PID      MemAlloc  StackBase/Ptr  Process
-----
1277      120632    7ffffcd0/7ffffefe4  xbar_client
1278       56800    7ffffce0/7fffffb5c  wwn
1279     1210220    7ffffce0/7fffffbac  vsan
1293      386144    7ffffcf0/7ffffebd4  span
1294     1396892    7ffffce0/7ffffdff4  snmpd
1295      214528    7ffffcf0/7ffff904   rscn
1296       42064    7ffffce0/7fffffb5c  qos
```

Where:

- MemAlloc = total memory allocated by the process.
- StackBase/Ptr = process stack base and current stack pointer in hex format

Displaying System Status

Use the **show system** command to display system-related status information (Example 31-7 to Example 31-10).

Example 31-7 Displays Default Switch Port States

```
switch# show system default switchport
System default port state is down
System default trunk mode is on
```

Example 31-8 Displays Error Information for a Specified ID

```
switch# show system error-id 0x401D0019
Error Facility: module
Error Description: Failed to stop Linecard Async Notification.
```

Example 31-9 Displays the System Reset Information

```
switch# Show system reset-reason module 5
----- reset reason for module 5 -----
1) At 224801 usecs after Fri Nov 21 16:36:40 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
2) At 922828 usecs after Fri Nov 21 16:02:48 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
3) At 318034 usecs after Fri Nov 21 14:03:36 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
4) At 255842 usecs after Wed Nov 19 00:07:49 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
```

The **show system reset-reason** command displays the following information:

- In a Cisco MDS 9500 Series switch, the last four reset-reason codes for the supervisor module in slot #5 and slot #6 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9200 Series switch, the last four reset-reason codes for supervisor module in slot #1 are displayed.
- The **show system reset-reason module number** command displays the last four reset-reason codes for a specific module in a given slot. If a module is absent, then the reset-reason codes for that module will not be displayed.

Use the **clear system reset-reason** command to clear the reset-reason information stored in NVRAM and volatile persistent storage. Use this command as listed below:

- In a Cisco MDS 9500 Series switch, this command clears the reset-reason information stored in NVRAM and volatile persistent storage in the active and standby supervisor modules.
- In a Cisco MDS 9200 Series switch, this command clears the reset-reason information stored in NVRAM and volatile persistent storage in the active supervisor module.

Example 31-10 Displays System Uptime

```
switch# show system uptime
Start Time: Sun Oct 13 18:09:23 2030
Up Time:    0 days, 9 hours, 46 minutes, 26 seconds
```

Use the **show system resources** command to display system-related CPU and memory statistics (see [Example 31-11](#)).

Example 31-11 Displays System-Related CPU and Memory Information

```
switch# show system resources
Load average:  1 minute: 0.43   5 minutes: 0.17   15 minutes: 0.11
Processes   :   100 total, 2 running
CPU states  :   0.0% user,   0.0% kernel,  100.0% idle
Memory usage: 1027628K total,   313424K used,   714204K free
               3620K buffers,   22278K cache
```

Where:

- Load is defined as number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states shows the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage provides the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the *used* memory statistics.

Configuring Core and Log Files

You can save cores (from the active supervisor module, the standby supervisor module, or any switching module) to an external flash (slot 0) or to a TFTP server in one of two ways:

- On demand—to copy a single file based on the provided process ID.
- Periodically—to copy core files periodically as configured by the user.

To copy the core and log files on demand, follow this step:

	Command	Purpose
Step 1	switch# copy core:7407 slot0:coreSample	Copies the core file with the process ID 7407 as coreSample in slot 0.
	switch# copy core://5/1524 tftp://1.1.1.1/abcd	Copies cores (if any) of a process with pid 1524 generated on slot 5 to tftp server.

- If the core file for the specified process ID is not available, you will see the following response:

```
switch# copy core:133 slot0:foo
No core file found with pid 133
```

- If two core files exist with same process ID, only one file will be copied:

```
switch# copy core:7407 slot0:foo1
2 core files found with pid 7407
Only "/isan/tmp/logs/calc_server_log.7407.tar.gz" will be copied to the destination.
```


To copy the core and log files periodically, follow these steps:

	Command	Purpose
Step 1	switch# config t	Enters configuration mode.
Step 2	switch(config)# system cores slot0:coreSample	Copies the core files coreSample to slot 0.
	switch(config)# system cores tftp://1.1.1.1/abcd	Copies the core files (abcd) in the specified directory on the TFTP server.
	switch(config)# no system cores	Disable the core files copying feature.

A new scheme overwrites any previously-issued scheme. For example, if you issue a new system core command, the cores are periodically saved to the new location or file.



Tip

Be sure to create any required directory before issuing this command. If the directory specified by this command does not exist, the switch software logs a syslog message each time a copy cores is attempted.)

Clearing the Core Directory

Use the **clear cores** command to clean out the core directory. The software keeps the last few cores per service and per slot and clears all other cores present on the active supervisor module.

```
switch# clear cores
```

Displaying Cores Status

Use the **show system cores** command to display the currently configured scheme for copying cores. See Examples 31-12 to 31-14.

Example 31-12 Displays the status of System Cores

```
switch# show system cores
Transfer of cores is enabled
```

Example 31-13 Displays All Cores Available for Upload from the Active Supervisor Module

```
switch# show cores
Module-num Process-name PID Core-create-time
-----
5 fspf 1524 Nov 9 03:11
6 fcc 919 Nov 9 03:09
8 acltcam 285 Nov 9 03:09
8 fib 283 Nov 9 03:08
```

Where:

module-num shows the slot number on which the core was generated. In this example, the `fsfp` core was generated on the active supervisor module (slot 5), `fcc` was generated on the standby supervisor module (slot 6), and `acltcam` and `fib` were generated on the switching module (slot 8).

Example 31-14 Displays Logs on the Local System

```
switch# show processes log
```

Process	PID	Normal-exit	Stack	Core	Log-create-time
ExceptionLog	2862	N	Y	N	Wed Aug 6 15:08:34 2003
acl	2299	N	Y	N	Tue Oct 28 02:50:01 2003
bios_daemon	2227	N	Y	N	Mon Sep 29 15:30:51 2003
capability	2373	N	Y	N	Tue Aug 19 13:30:02 2003
core-client	2262	N	Y	N	Mon Sep 29 15:30:51 2003
fcanalyzer	5623	N	Y	N	Fri Sep 26 20:45:09 2003
fcd	12996	N	Y	N	Fri Oct 17 20:35:01 2003
fcdomain	2410	N	Y	N	Thu Jun 12 09:30:58 2003
ficon	2708	N	Y	N	Wed Nov 12 18:34:02 2003
ficonstat	9640	N	Y	N	Tue Sep 30 22:55:03 2003
flogi	1300	N	Y	N	Fri Jun 20 08:52:33 2003
idehsd	2176	N	Y	N	Tue Jun 24 05:10:56 2003
lmgrd	2220	N	N	N	Mon Sep 29 15:30:51 2003
platform	2840	N	Y	N	Sat Oct 11 18:29:42 2003
port-security	3098	N	Y	N	Sun Sep 14 22:10:28 2003
port	11818	N	Y	N	Mon Nov 17 23:13:37 2003
rlir	3195	N	Y	N	Fri Jun 27 18:01:05 2003
rscn	2319	N	Y	N	Mon Sep 29 21:19:14 2003
securityd	2239	N	N	N	Thu Oct 16 18:51:39 2003
snmpd	2364	N	Y	N	Mon Nov 17 23:19:39 2003
span	2220	N	Y	N	Mon Sep 29 21:19:13 2003
syslogd	2076	N	Y	N	Sat Oct 11 18:29:40 2003
tcap	2864	N	Y	N	Wed Aug 6 15:09:04 2003
tftpd	2021	N	Y	N	Mon Sep 29 15:30:51 2003
vpm	2930	N	N	N	Mon Nov 17 19:14:33 2003

Configuring Kernel Core Dumps

**Caution**

Changes to the kernel cores should be made by an administrator or individual who is completely familiar with switch operations.

When a specific module's operating system (OS) crashes, it is sometimes useful to obtain a full copy of the memory image (called a kernel core dump) to identify the cause of the crash. When the module experiences a kernel core dump it triggers the proxy server configured on the supervisor. The supervisor sends the module's OS kernel core dump to the Cisco MDS 9000 System Debug Server. Similarly, if the supervisor OS fails the supervisor sends its OS kernel core dump to the Cisco MDS 9000 System Debug Server.

**Note**

The Cisco MDS 9000 System Debug Server is a Cisco application that runs on Linux. It creates a repository for kernel core dumps. You can download the Cisco MDS 9000 System Debug Server from the Cisco.com website at <http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml>.

Kernel core dumps are only useful to your technical support representative. The kernel core dump file, which is a large binary file, must be transferred to an external server that resides on the same physical LAN as the switch. The core dump is subsequently interpreted by technical personnel who have access to source code and detailed memory maps.

**Tip**

Core dumps take up disk space on the Cisco MDS 9000 System Debug Server application. If all levels of core dumps (**level all** option) are configured, you need to ensure that a minimum of 1GB of disk space is available on the Linux server running the Cisco MDS 9000 System Debug Server application to accept the dump. If the process does not have sufficient space to complete the generation, the module resets itself.

To configure the external server, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# kernel core target 10.50.5.5 succeeded	Configures the external server's IP address.

To configure the module information, follow these steps:

	Command	Purpose
Step 1	switch# config terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# kernel core module 5 succeeded	Configures kernel core generation for module 5.
	switch(config)# kernel core module 5 level header succeeded	Configures kernel core generation for module 5, and limits the generation to header-level cores.
Step 3	switch(config)# kernel core limit 2 succeeded	Configures generations for two modules. The default is 1 module.

All changes made to kernel cores are saved to the running configuration and may be viewed using the **show running-config** command. Alternatively, use the **show kernel cores** command to view specific configuration changes (see examples 31-15 to 31-17).

Example 31-15 Displays the Core Limit

```
switch# show kernel core limit
2
```

Example 31-16 Displays the External Server

```
switch# show kernel core target
10.50.5.5
```

Example 31-17 Displays the Core Settings for the Specified Module

```
switch# show kernel core module 5
module 5 core is enabled
  level is header
  dst_ip is 10.50.5.5
  src_port is 6671
  dst_port is 6666
  dump_dev_name is eth1
  dst_mac_addr is 00:00:0C:07:AC:01
```




Symbols

*(wildcard)

port security authorization [18-5](#)

Numerics

16-port modules

BB_credits [10-11](#)

LEDs [10-13](#)

preserving configurations [7-6](#)

See also switching modules

32-port modules

BB_credits [10-11](#)

configuration guidelines [10-8, 12-2](#)

preserving configurations [7-6](#)

SPAN guidelines [28-6](#)

See also switching modules

A

AAA

authorization and authentication process [16-16](#)

setting authentication [16-15](#)

usage [1-10, 16-1](#)

user accountability [16-24](#)

access control

iSCSI enforcement [22-60](#)

iSCSI mechanisms [22-58](#)

Access Control Lists

See ACLs

access-group configuration [20-9](#)

accounting [16-24](#)

ACLs

adding entries [20-6](#)

clearing counters [20-10](#)

configuration guidelines [20-5](#)

creating [20-5](#)

defining [20-6](#)

log dump [20-10](#)

operands [20-7](#)

removing entries [20-8](#)

activating

IVR topology [14-7](#)

active

modules [7-2](#)

states [5-6](#)

zones [13-7](#)

zone sets [13-11](#)

adding

IP addresses [20-20](#)

ports to a PortChannel [12-6](#)

SNMP communities [16-34](#)

switches [4-4](#)

adding ACL entries [20-6](#)

address-allocation cache [24-15](#)

Address Resolution Protocol

See ARP

administrative speed

configuring [10-10](#)

administrative states

description [10-6](#)

administrator passwords

configuring [4-9](#)

configuring switch [4-3](#)

- creating additional accounts [4-5](#)
- default [4-5](#)
- recovering [16-26](#)
- requirements (note) [4-6, 4-10](#)
- advertisement packets
 - setting time intervals [20-21](#)
- aggregated flow statistics [19-14](#)
- aliases
 - configuring [13-5](#)
- ALPA [10-27](#)
- applying ACLs [20-8](#)
- area FCID
 - configuring [24-11](#)
- ARP
 - clearing and viewing entries [20-12](#)
 - IP services [1-6](#)
- ARP caches
 - clearing [22-8](#)
 - managing [22-8](#)
- assigning
 - alias names [13-5](#)
 - contact information [23-2](#)
 - domain IDs [24-4](#)
 - FC IDs [19-9](#)
 - global keys [16-7](#)
 - host key [16-6](#)
 - users [16-34](#)
 - zone members [13-4](#)
- authentication
 - CHAP option [22-79](#)
 - global override [22-61](#)
 - iSCSI setup [22-78](#)
 - See MD5 authentication
 - See simple text authentication
- authentication, authorization, and accounting
 - See AAA
- automatic synchronization
 - conditions [5-6](#)
- autonomous fabric ID [14-6](#)

- AutoNotify
 - destination profile (note) [23-5](#)
 - registration [23-3](#)
 - service contract [23-2](#)
- auto port mode
 - configuring [10-10](#)
 - description [10-5](#)
 - interface configuration [10-2](#)

B

- basic input/output system
 - See BIOS
- BB_credits
 - configuring [10-11](#)
 - reason codes [10-7](#)
- beacon mode
 - configuring [10-13](#)
 - identifying LEDs [10-13](#)
 - LEDs [7-10](#)
- Berkeley Packet Filter
 - See BPF
- BIOS
 - boot sequence [6-26](#)
 - recovering corrupted bootflash [6-27](#)
 - recovery sequence [6-27](#)
 - setup (figure) [6-29](#)
- BIOS upgrades [6-21](#)
- boot
 - sequence [6-26](#)
 - variable synchronization [5-4](#)
- bootflash
 - copying to [6-17](#)
 - description [2-17](#)
 - device [6-17](#)
 - file system [6-1](#)
 - initializing [2-18](#)
 - recovering corrupted [6-26 to 6-27](#)
 - space requirements [6-2](#)

- See also internal bootflash
- bootloader
 - nondisruptive upgrades [6-19](#)
 - skipping phases [6-31](#)
- bootup diagnostics [7-3](#)
- boot variables
 - disruptive upgrades [6-23](#)
- border switch [14-3](#)
- BPF
 - library [29-16](#)
 - See also libpcap freeware
- B ports
 - functionality [10-5](#)
- broadcast
 - in-band addresses default [7-12](#)
 - routing [19-10](#)
- buffer-to-buffer credits
 - See BB_credits
- build fabric frames [24-3](#)

C

- cache
 - See address-allocation cache
- Call Home
 - configuring [23-3 to 23-9](#)
- Call Home
 - functionality [1-7](#)
 - message format options [23-2](#)
- capture filters [29-16](#)
- CDP
 - clearing [4-38](#)
 - configuring [4-37](#)
 - configuring globally [4-38](#)
 - displaying [4-39](#)
 - hold time [4-38](#)
 - packet transmission [4-37](#)
- CHAP authentication [22-79](#)
- chassis
 - types [7-1](#)
- checks
 - See compatibility checks
- Cisco Discovery Protocol
 - See CDP
- Cisco MDS 9000 System Debug Server [31-8](#)
- Cisco MDS 9120 Directors
 - overview [1-1](#)
- Cisco MDS 9140 Directors
 - overview [1-1](#)
- Cisco MDS 9200 Series
 - LEDs [7-8](#)
 - mgmt0 LEDs [7-9](#)
 - supervisor modules [7-2](#)
- Cisco MDS 9216 switches
 - high availability [1-4, 5-2](#)
 - modules [1-9, 7-1](#)
 - overview [1-1](#)
 - supervisor module [7-4](#)
- Cisco MDS 9500 Series
 - high availability [1-4, 5-2](#)
 - LEDs [7-9](#)
 - overview [1-2](#)
 - supervisor modules [7-2](#)
- Cisco MDS 9506 Directors
 - modules [1-9, 7-1](#)
 - overview [1-1, 1-2](#)
- Cisco MDS 9509 Directors
 - modules [1-9, 7-1](#)
 - overview [1-1, 1-2](#)
- claim certificate [3-2](#)
- clearing
 - FIB statistics [19-15](#)
 - FSPF counters [19-9](#)
 - IVZ database [14-12](#)
 - zone sets [13-12](#)
- CLI
 - accessing submodes [2-3](#)
 - alternative [1-11](#)

- command modes [2-3](#)
- updating SNMPv3 passwords [16-33](#)
- clock modules
 - monitoring status [8-10](#)
- CMOS
 - configuration [6-29](#)
 - saving changes [6-30](#)
- COM 1
 - configuring [4-32](#)
- command-line interface
 - See CLI
- commands
 - saving output to files [2-22](#)
- CompactFlash
 - devices [2-17, 2-18, 6-17](#)
 - disk [6-1](#)
 - slot 0 [6-17](#)
- compatibility checks [12-7](#)
- computing routes [19-1](#)
- configuring
 - FCIP links [22-20](#)
 - IVR [14-5](#)
 - unique area FCIDs [24-11](#)
- congestion control methods
 - See edge quench congestion control
 - See FCC
- connecting a modem
 - COM 1 [4-33](#)
 - console [4-33](#)
- consistent switch states [12-6](#)
- console session
 - severity levels [26-5](#)
- control frames [22-18](#)
- control traffic
 - disabling [25-4](#)
- core dumps
 - IPS module [22-16](#)
 - kernel [31-8](#)
- cores [31-6](#)

- counted licenses [3-2](#)
- creating
 - IVR topology [14-6](#)
- creating SNMP roles [16-32](#)
- cross-VSAN communication [14-8](#)
- customized
 - targets [27-4](#)

D

- databases
 - See zone databases
- data field
 - configuring size [10-12](#)
- data frames [22-18](#)
- dead time interval [19-7](#)
- default gateway
 - BIOS setup configuration [6-29](#)
 - configuring mgmt0 Ethernet interfaces [10-16](#)
 - recovering loader> prompt [6-31](#)
 - recovering switch(boot)# prompt [6-32](#)
- default groups [16-33](#)
- default zones
 - description [13-9](#)
 - interoperability [29-21](#)
- deleting
 - FSPF configurations [19-5](#)
 - PortChannels [12-6](#)
- deny conditions [20-5](#)
- destination IDs
 - exchange based [12-5](#)
 - flow based [12-4](#)
 - frame identification [25-2](#)
 - frame loop back [29-4](#)
 - in-order delivery [19-10, 25-2](#)
 - load balancing [1-6, 12-1](#)
 - path selection [9-6](#)
- destination profiles
 - configuring [23-5](#)

- device IDs
 - Call Home format [23-16, 23-17](#)
 - copying files [16-33](#)
 - report capacity [27-1](#)
- Device View
 - description [1-11](#)
- digital signature algorithm
 - See DSA key pairs
- Dijkstra's algorithm [19-2](#)
- direct memory access
 - See DMA-bridge
- disabling routing protocols [19-5](#)
- discovered
 - LUNs [27-4](#)
 - targets [27-3](#)
- display filters
 - selective viewing [29-12](#)
- disruptive
 - upgrades [6-4](#)
- distribution tree [19-10](#)
- DMA-bridge [22-10](#)
- documentation
 - related documents [xxxvi](#)
- domain ID
 - IVR guidelines [14-4](#)
- domain IDs
 - configuring [24-4](#)
 - distributing [24-2](#)
 - failure [10-7](#)
 - interoperability [29-21](#)
 - preferred [24-5](#)
 - range [2-25](#)
 - static [24-5](#)
 - unique [14-5](#)
- domain manager
 - isolation [10-7](#)
- domain names
 - defining [20-24](#)
- Domain Name System servers

- See DNS servers
- domain overlap
 - isolation [10-7](#)
- drop latency time
 - configuring [19-13](#)
- dsa key pairs
 - generating [16-27](#)

E

- edge
 - switch [14-3](#)
 - VSAN [14-3](#)
- edge quench congestion control
 - description [25-2](#)
- egress port [28-12, 28-25](#)
- EISL
 - functionality [1-6](#)
 - PortChannel links [12-1](#)
- ELP failure [10-7](#)
- e-mail notification
 - Call Home [23-1](#)
- enabling
 - IVR [14-6](#)
- enforcing licenses [3-1](#)
- Enterprise package [3-4](#)
- environmental monitors [7-8, 7-9](#)
- E ports
 - 32-port guidelines [10-8, 12-2](#)
 - classes of service [10-3](#)
 - configuring [10-10](#)
 - FSPF topology [19-2](#)
 - interface modes [10-2](#)
 - isolation [10-7](#)
 - recovering from isolation [13-10](#)
 - SPAN [28-3](#)
 - trunking [1-6](#)
 - trunking configuration [11-3](#)
- error disabled code [10-7](#)

error messages
 description [26-2](#)

error state [6-34](#)

ESC failure [10-7](#)

Ethereal freeware
 analyzer [29-8](#)
 information [29-7](#)

Ethernet PortChannels
 configuring [22-14](#)

evaluation license [3-3](#)

exchange IDs
 in-order delivery [19-10](#)
 load balancing [1-6, 12-1, 29-4](#)
 path selection [9-6](#)

exchange link parameter
 See ELP failure

expiry alerts
 licenses [3-10](#)

exporting
 zone databases [13-10](#)

extended ISL
 See EISL

external RADIUS server
 CHAP [22-79](#)

external server
 configuring [31-9](#)

F

fabric
 See build fabric frames
 See reconfigure fabric frames

Fabric Analyzer
 capture range [2-25](#)
 configuring [29-9](#)
 description [29-7](#)
 frame range [2-25](#)

Fabric Configuration Server
 See FCS

fabric login
 See FLOGI

Fabric Manager
 description [1-9](#)
 Device View [1-11](#)
 Fabric View [1-11](#)

fabric names
 setting [24-8](#)

fabric pWWNs
 configuring zones [13-4](#)
 zone membership [13-2](#)

fabric reconfiguration
 fcdomain phase [24-2](#)

fabric shortest path first
 See FSPF

Fabric View
 description [1-11](#)

fail over protection [22-12](#)

fan modules
 monitoring status [8-10](#)

fan trays
 overview [1-2](#)

fault tolerant fabric
 example (figure) [19-2](#)

FC aliases
 configuring zones [13-4](#)

fcanalyzer
 clearing hosts [29-11](#)
 displaying filters [29-12](#)

FCC
 benefits [25-2](#)
 default settings [25-12](#)
 enabling [25-3](#)
 frame handling [25-2](#)
 logging facility [26-2](#)

fcdomain
 configuring [24-1](#)
 default settings [24-16](#)

FC IDs

- address format [2-25](#)
- allocating [24-2, 29-19](#)
- allocating areas [29-19](#)
- configuring zones [13-4](#)
- FCIP
 - configuring [22-17](#)
 - Gigabit Ethernet ports [22-4](#)
 - interfaces [22-18, 22-19](#)
 - IPS module [22-2, 22-17](#)
 - IP storage [22-1](#)
 - links [22-18](#)
 - profiles [22-18, 22-19](#)
 - TCP connections [22-18](#)
- FCIP parameters
 - default [22-95](#)
- Fcot not present [10-7](#)
- fcping
 - invoking [29-5](#)
- FCS
 - configuring [30-3](#)
 - description [30-2](#)
 - logging facility [26-2](#)
 - significance [30-3](#)
- fctrace
 - invoking [29-4](#)
- feature-based licensing [3-4](#)
- Fibre Channel analyzers [28-10](#)
- Fibre Channel Congestion Control
 - See FCC
- Fibre Channel domain
 - See fcdomain
- Fibre Channel over IP
 - See FCIP
- Fibre Channel traffic
 - SPAN sources [28-3](#)
- file system
 - formatting [2-18](#)
 - redirection [2-22](#)
 - volatile [2-19](#)
- File Transfer Protocol
 - See FTP
- filters
 - capture [29-16](#)
 - defining display [29-13](#)
- FLOGI
 - displaying details [15-1](#)
 - logging facility [26-2](#)
- flow statistics [19-14](#)
- FL ports
 - classes of service [10-4](#)
 - configuring [10-10](#)
 - fctrace [29-4](#)
 - interface modes [10-2](#)
 - nonparticipating code [10-8](#)
 - persistent FC IDs [24-10](#)
 - SPAN [28-3](#)
- FM Server package [3-5](#)
- F ports
 - classes of service [10-3](#)
 - configuring [10-10](#)
 - interface modes [10-2](#)
 - SPAN [28-3](#)
- frames
 - control [22-18](#)
 - data [22-18](#)
 - encapsulation [10-12, 28-8](#)
 - flow [1-8](#)
 - MTU [22-5](#)
 - reordering [19-10](#)
- FSPF
 - alternative paths [19-1](#)
 - clearing counters [19-9](#)
 - computing link cost [19-6](#)
 - configuring globally [19-4](#)
 - configuring on interfaces [19-6](#)
 - default settings [19-20](#)
 - disabling on interfaces [19-7](#)
 - disabling routing protocols [19-5](#)

- hold time range [2-25, 19-1](#)
- interoperability [29-22](#)
- link state protocol [19-2](#)
- reconvergence time [19-2](#)
- routing services [19-1](#)
- topologies example [19-2](#)

FTP

- logging facility [26-2, 26-7](#)

full core dumps

- IPS modules [22-16](#)

full zone set

- considerations [13-7](#)
- distribution [13-11](#)

Fx ports

- 32-port default [10-8](#)
- configuring [10-10](#)
- FCS [30-2](#)
- interface modes [10-5](#)

G

Gigabit Ethernet

- configuring [22-5](#)
- IP routing [22-7](#)
- major interfaces [22-6](#)
- ports [22-4](#)
- subinterfaces [22-6](#)

Gigabit Ethernet parameters

- default [22-95](#)

global iSCSI information

- displaying [22-67](#)
- grace period [3-3](#)

H

hardware

- displaying inventory [8-2](#)
- status description [7-3](#)

- hard zoning [13-9](#)

- HA-standby [5-4, 7-3](#)

HBA ports

- configuring area FCIDs [24-11](#)

- hello time interval [19-6](#)

- hidden routes [19-16](#)

high availability

- Ethernet PortChannel [22-78](#)
- features [22-75](#)
- functionality [1-4, 5-2](#)
- licensing [3-5](#)
- process restartability [5-4](#)
- software upgrade [6-4](#)
- status [5-5](#)
- VRRP [22-77](#)
- VRRP features [22-12](#)

host ID

- licensing [3-2](#)

I

ICMP packets

- type value [20-8](#)

ICMP statistics

- displaying [22-12](#)

identical passwords

- CLI and SNMP [16-33](#)

IDs

- CCO IDs [23-3](#)
- contract IDs [23-4, 23-16](#)
- customer IDs [23-4](#)
- image version and IDs [6-1](#)
- login IDs [4-5](#)
- process IDs [4-27, 31-2, 31-6](#)
- profile IDs [23-5](#)
- region ID [19-4](#)
- serial IDs [23-17](#)
- server IDs [23-17](#)
- site IDs [23-4, 23-16](#)

- See destination IDs
- See device IDs
- See domain IDs
- See exchange IDs
- See FC IDs
- See port IDs
- See source IDs
- See user IDs
- See VR IDs
- See VSAN IDs
- images
 - See kickstart images
 - See software images
 - See system images
- software upgrades
 - See also one-step upgrade
 - See also step-by-step upgrades
- importing database [13-10](#)
- inactive code [10-7](#)
- inconsistent switch states [12-6](#)
- incremental licenses [3-2](#)
- ingress port [28-11](#)
- initiator access list [22-59](#)
- in-order delivery [19-10](#)
 - enabling [19-13](#)
- in-order guarantee [19-11](#)
- install all
 - command benefits [6-6](#)
 - command examples [6-10](#)
 - command failure cases [6-7](#)
 - command function [6-6](#)
 - command requirements [6-3](#)
 - command usage [6-7](#)
 - remote location path (caution) [6-12](#)
- installing
 - licenses [3-5](#)
- insufficient power [7-3](#)
- interface
 - adding to PortChannels [12-7](#)
 - configuring FSPF [19-6](#)
 - suspended states [12-7](#)
- interfaces
 - characteristics [10-2](#)
 - configuring [10-9](#)
 - data field size [10-12](#)
 - default settings [10-15](#)
 - description [10-10](#)
 - modes [10-2, 10-9](#)
 - reason codes [10-6](#)
 - states [10-6](#)
- internal bootflash
 - description [2-17](#)
 - flash devices [2-17](#)
 - See also bootflash
- internal switch states
 - description [5-7](#)
- interoperability
 - configuring [29-21](#)
 - verifying status [29-23 to ??](#)
- inter-switch links
 - See ISL
- Inter-VSAN Routing
 - see IVR
- invoking fcping [29-5](#)
- IP Access Control Lists
 - See ACLs
- IP address
 - address format [2-25](#)
 - SMTP server [23-9](#)
- IP addresses
 - configuring in VSANs [20-11](#)
- IPFC
 - logging facility [26-3](#)
- IP features
 - default settings [20-25](#)
- IP forwarding
 - disabling [20-11](#)
- IP over Fibre Channel

- See IPFC
- IP routing
 - Gigabit Ethernet interface [22-7](#)
 - static [1-6](#)
- IPS core dumps
 - See core dumps
- IP services
 - default settings [20-24](#)
- IPS module
 - CDP [4-37](#)
 - CDP support [22-16](#)
 - functions [22-2](#)
 - port mode [22-4](#)
- IPS modules
 - core dumps [22-16](#)
- IPS Ports
 - multiple connections [22-76](#)
- IPS ports [22-2](#)
- IPS services module
 - See IPS module
- IPS statistics
 - displaying [22-73](#)
- IP statistics
 - displaying [22-10](#)
- IP storage
 - VRRP [22-13](#)
- iSCSI
 - Gigabit Ethernet ports [22-4](#)
 - IPS module [22-2](#)
- iSCSI authentication [22-60](#)
- iSCSI discovery [22-60](#)
- iSCSI host
 - multiple VSANs [22-55](#)
- iSCSI hosts
 - configuring accessibility [22-59](#)
- iSCSI initiators
 - assigning WWNs [22-54](#)
 - displaying [22-69](#)
 - dynamic mapping [22-53](#)
 - identifying [22-53](#)
- iSCSI interfaces
 - displaying [22-64](#)
- iSCSI parameters
 - default [22-95](#)
- iSCSI session creation [22-60](#)
- iSCSI sessions
 - displaying [22-67](#)
- iSCSI targets
 - access control [22-58](#)
 - secondary access [22-49](#)
- iSCSI user information
 - displaying [22-75](#)
- iSCSI virtual targets
 - displaying [22-73](#)
- ISL
 - PortChannel links [12-1](#)
- isolation
 - reason codes [10-7](#)
- IVR
 - path [14-3](#)
 - sharing resources [14-2](#)
 - zone [14-3](#)
 - zonesets [14-3](#)

J

- jumbo frames
 - see MTU frame size

K

- kernel core dumps [31-8](#)
 - configuring [31-9](#)
- kickstart images
 - downloading [2-18](#)
 - KICKSTART variable [6-2](#)
 - loading system images [6-26](#)

- overview [6-1](#)
- recovering corrupted [6-31](#)
- recovery [6-32](#)
- recovery interruption [6-27](#)

L

LEDs

- identifying beacon [10-13](#)

- libpcap freeware [29-7](#)

license

- terminology [3-2](#)

licenses

- factory-installed [3-6](#)

licensing

- installing [1-4, 3-1](#)

- link cost [19-2](#)

- link end points [22-18](#)

- link failure [10-7](#)

- high availability [5-2](#)

- link redundancy

- Ethernet PortChannels [22-14](#)

- load balancing [12-1](#)

- attributes [9-6](#)

- guarantee [9-7](#)

- mechanisms [12-4](#)

loader

- loading kickstart [6-26](#)

- local capture [29-9](#)

- log files [31-6](#)

- configuring [26-6](#)

logging

- default settings [26-12](#)

- severity levels [26-4](#)

- system messages [26-2](#)

- logging levels

- IVR [14-13](#)

- logical unit numbers

- See LUNs

- loop monitoring [29-20](#)

- loop port [29-20](#)

- LSR [19-18](#)

LUNs

- address format [2-25](#)

- displaying discovered, example [27-4](#)

M

- MAC= keyword [20-10](#)

MAC address

- format [2-25](#)

- Mainframe package [3-5](#)

- major threshold [7-8](#)

Management Information Base

- See MIB

- management module [7-8](#)

- management redundancy

- high availability [5-2](#)

- manual assignment [22-53](#)

- mapping

- iSCSI hosts [22-52](#)

- MD5 authentication [20-22](#)

membership

- IVR zones [14-8](#)

- memory test [6-27, 6-28](#)

mgmt0 interfaces

- autosensing port [10-16](#)

- configuring [10-16](#)

- configuring ethernet ports [20-3](#)

- overview [10-1](#)

- recovery from switch(boot)# prompt [6-32](#)

- minor threshold [7-8](#)

- modify existing users [16-33](#)

module

- configuring logging [26-5](#)

- module-based licensing [3-4](#)

- module configuration

- sample scenarios [7-6](#)

module status [10-1](#)
 module temperature [8-9](#)
 monitoring traffic [28-7, 28-17](#)
 MTU frame size [22-5](#)
 multicast routing [19-10](#)
 multi-pid option [15-9](#)

N

name server
 interoperability [29-22](#)
 name server proxy [15-3](#)
 native VSAN [14-3](#)
 network administrator [2-24, 16-3](#)
 network operator [2-24, 16-3](#)
 Network Time Protocol
 See NTP
 network traffic
 monitoring [28-7, 28-17](#)
 next hop domain ID [19-9](#)
 NL ports
 fctrace [29-4](#)
 interface modes [10-5](#)
 zone enforcement [13-9](#)
 node-locked license [3-2](#)
 node WWNs
 See nWWNs
 nondisruptive
 restart [5-2](#)
 upgrades [6-4](#)
 nonparticipating code [10-8](#)
 non-trunking ISL [11-2](#)
 nonvolatile storage [7-6](#)
 N ports
 fctrace [29-4](#)
 zone enforcement [13-9](#)
 zone membership [13-2](#)
 NTP
 logging facility [26-3](#)

nWWNs
 address format [2-25](#)
 Nx ports
 hard zoning [13-9](#)

O

offline code [10-7](#)
 one-step upgrade
 install all command [6-4](#)
 reload command [6-4](#)
 operational interfaces
 viewing PortChannels [12-10](#)
 operational state [10-10](#)
 operational state setting
 description [10-6](#)
 originator exchange IDs
 See exchange IDs
 out-of-order delivery [19-10](#)

P

partial core dumps
 IPS module [22-16](#)
 password recovery [16-26](#)
 path discovery [29-4](#)
 permanent license [3-3](#)
 permit conditions [20-5](#)
 permitted filters [29-17](#)
 persistent FC ID [24-9](#)
 persistent FC IDs
 displaying [24-14](#)
 physical interfaces [22-10](#)
 policy
 IVR zones [14-8](#)
 port aggregation [5-2](#)
 PortChannel
 configuring FC routes [19-9](#)

- functionality [1-6](#)
- high availability [5-2](#)
- in-order guarantee [19-11](#)
- link changes [19-11](#)
- link failure [19-3](#)
- load balancing [1-6](#)
- logging facility [26-3](#)
- membership [12-8](#)
- range [2-25](#)
- reason codes [10-8](#)
- PortChannels
 - adding Gigabit Ethernet interfaces [22-16](#)
 - adding interfaces [12-6](#)
 - configuring [12-5](#)
 - default settings [12-12](#)
 - deleting [12-6](#)
 - examples [12-2](#)
 - forcing additions [12-7](#)
 - guidelines [12-9](#)
 - interoperability [29-22](#)
 - member combinations [22-14](#)
 - SPAN [28-3](#)
 - trunking comparison [12-3](#)
- port group [10-8, 12-2](#)
- port IDs
 - configuring zones [13-4](#)
- port mode
 - IPS [22-2, 22-4](#)
- port modes
 - auto [10-5](#)
- ports
 - virtual E [22-18](#)
- Port world wide name
 - See pWWN
- port WWNs
 - See pWWNs
- power supplies [1-2, 1-3, 7-8, 7-9](#)
 - configuring [8-6](#)
 - displaying configuration [8-6](#)
 - guidelines [8-6](#)
 - modes [4-29, 8-6](#)
- power usage
 - displaying details [8-5](#)
- preempt option [20-21](#)
- preferred domain IDs [24-5](#)
- preshared key [16-7](#)
- principal switch [24-4, 24-5](#)
 - selecting [24-1](#)
- private device [10-27](#)
- process ID [31-6](#)
- Process Logs [31-4](#)
- process restartability [5-4](#)
- product authorization key [3-2](#)
- proof of purchase [3-2](#)
- protocol analysis [29-7](#)
- pWWNs
 - address format [2-25](#)
 - configuring zones [13-4](#)
 - zone membership [13-2](#)

Q

- QoS
 - default settings [25-12](#)
 - displaying information [25-4, 25-9](#)
 - enabling [25-6](#)
 - enabling control traffic [25-4](#)
 - logging facilities [26-3](#)
 - priority queuing [1-8](#)
- quality of service
 - See QoS

R

- R_A_TOV time [10-7](#)
- RADIUS
 - AAA solutions [1-10, 16-1](#)

- configured parameters [16-9](#)
- secret key [1-10, 16-1](#)
- setting preshared key [16-7](#)
- specifying servers [16-6](#)
- specifying time-out [16-8](#)
- rebooting switch [7-5](#)
- reconfigure fabric [10-7](#)
- reconfigure fabric frames [24-3](#)
- reconvergence time
 - FSPF [19-2](#)
- recovering passwords [16-26](#)
- recovery sequence [6-27](#)
- redundancy states [5-6](#)
- redundant physical links [19-3](#)
- Registered State Change Notification
 - See RSCN
- remote capture [29-9, 29-11](#)
- remote capture daemon [29-8](#)
- Remote Capture Protocol
 - See RPCAP
- Remote Monitoring
 - See RMON
- retransmit intervals [19-8](#)
- roles
 - additional [16-3](#)
- route cost
 - computing [19-6](#)
- route table [22-7](#)
- routing
 - See broadcast routing
 - See IP routing
- RPCAP
 - Ethereal communication [29-8](#)
- rsa1 key pairs
 - generating [16-27](#)
- rsa key pairs
 - generating [16-27](#)
- RSCN
 - logging facility [26-3](#)

- run time checks [19-8](#)

S

- SAN extension package [3-4](#)
- SAN operating system
 - See SAN-OS
- SAN-OS [6-2](#)
- SCSI LUNs
 - discovering targets [27-1](#)
- SD ports
 - bidirectional traffic [28-12](#)
 - configuring [10-10, 28-7, 28-22](#)
 - interface modes [10-2, 10-4](#)
- secondary MAC address [29-18](#)
- Secure Shell
 - See SSH
- security control
 - local [16-2, 16-15](#)
 - remote [16-2, 16-6, 16-10](#)
- security features
 - default settings [16-36, 17-10, 18-12](#)
- security parameter index
 - See SPI
- See MAC address
 - See also WWNs
- selective purging
 - persistent FC IDs [24-13](#)
- severity levels
 - logging [26-5](#)
- shutdown state [10-8, 12-2](#)
- Simple Network Management Protocol
 - See SNMP
- simple text authentication [20-22](#)
- simulating
 - Call Home [23-10](#)
- slot0
 - formatting [2-18](#)
- small computer system interface

- See SCSI
- SMARTnet [23-3](#)
- SMTP
 - server address [23-9](#)
- SNMP
 - access control [16-30](#)
 - access groups [16-31](#)
 - CLI configuration [16-30](#)
 - community strings [16-30](#)
 - configuring from CLI [16-33](#)
 - counter Information [16-35](#)
 - displaying information [16-35](#)
 - read-write access [16-34](#)
 - server contact [23-2](#)
 - versions [16-30](#)
- SNMP manager
 - FCS [30-3](#)
- SNMPv3
 - security features [16-30](#)
- software image
 - default setting [6-35](#)
 - error state [6-26](#)
 - recognizing errors [6-34](#)
- software images
 - bootflash corruption [6-26](#)
 - compatibility issues [6-19](#)
 - corruption [6-26](#)
 - recovery procedure [6-26](#)
 - space requirement [6-2](#)
 - synchronizing [5-4](#)
 - upgrade requirements [6-2](#)
 - upgrading [6-1](#)
 - variables [6-2](#)
- software upgrades
 - disruptive [6-7](#)
 - high availability [5-2](#)
 - manual, dual supervisor [6-17](#)
 - mechanisms [6-4](#)
 - quick [6-23](#)
- soft zoning [13-9](#)
- source IDs
 - Call Home event format [23-16](#)
 - exchange based [12-5](#)
 - flow based [12-4](#)
 - frame identification [25-2](#)
 - frame loop back [29-4](#)
 - in-order delivery [19-10](#)
 - load balancing [1-6, 12-1](#)
 - path selection [9-6](#)
- SPAN
 - configuring sessions [28-5](#)
 - default settings [28-14](#)
 - egress source [28-3](#)
 - encapsulating frames [28-8](#)
 - FC analyzers [28-10](#)
 - ingress source [28-2](#)
 - monitoring traffic [1-8, 28-2](#)
 - source configuration [28-4](#)
 - sources [28-4](#)
- speed
 - LEDs [7-10](#)
- SPI
 - configuring virtual router [20-22](#)
- SSH
 - default service [16-27](#)
 - force option [16-28](#)
 - host key pair [16-27](#)
 - protocol status [16-29](#)
 - session [6-17](#)
- SSH session
 - message logging [26-4](#)
- Standard package [3-4](#)
- standby module [7-2](#)
 - monitoring [5-2](#)
- standby supervisor [5-4](#)
- stateful
 - HA-switchover [5-3](#)
- static domain IDs [24-5](#)

- static mapping
 - iSCSI initiators [22-53](#)
 - pWWN assignment [22-54](#)
- static routes [9-8](#)
 - run time checks [19-8](#)
- status
 - LEDs [7-10](#)
- storage
 - permanent and temporary [2-17](#)
- storage devices
 - access control [13-1](#)
- ST ports
 - configuring [28-19](#)
- subnet mask
 - BIOS setup configuration [6-29](#)
 - configuring IP routes [20-12](#)
 - configuring mgmt0 [4-20](#)
 - configuring mgmt0 interfaces [10-16, 20-2](#)
 - configuring switch [4-3](#)
 - default setting [7-12](#)
 - initial configuration [4-6, 4-10](#)
 - loader> prompt recovery [6-31](#)
 - switch(boot)# prompt recovery [6-32](#)
- subnetwork requirements [22-6](#)
- subordinate switch [24-7](#)
- supervisor module
 - CDP support [22-16](#)
 - default settings [7-12](#)
- supervisor modules
 - active [1-9, 5-2](#)
 - active state [5-6, 5-7, 7-3](#)
 - default settings [7-12](#)
 - dual modules [7-2](#)
 - high availability [5-2](#)
 - major threshold [7-9](#)
 - recovering password [6-33](#)
 - resetting [7-5](#)
 - standby module [1-9](#)
 - standby state [5-7, 7-3](#)
 - standby status [7-3](#)
 - states [5-6](#)
 - switch options [1-9](#)
 - switchover [5-3](#)
 - viewing information [7-4](#)
- suspended state [12-7](#)
- switch
 - dual supervisor [6-33](#)
 - reliability service [1-4](#)
 - reloading [7-5](#)
 - role-based access [1-10](#)
 - secure access [1-10](#)
 - security management [1-9](#)
 - single supervisor [6-32](#)
 - SNMPv3 access [1-10](#)
 - verifying modules [7-2](#)
- switchability
 - high availability [5-2](#)
- switched port analyzer
 - See SPAN
- switching module
 - 16-port [7-6](#)
 - 32-port [7-6](#)
 - image [7-2](#)
 - LEDs [7-8](#)
 - power cycle [7-5](#)
 - powering off [7-7](#)
 - reloading [7-5](#)
 - status [7-2](#)
 - viewing states [7-3](#)
- switching modules
 - connecting to [7-4](#)
 - LEDs [7-8](#)
 - LEDs (table) [7-10](#)
 - managing [7-1](#)
 - powering off [7-7](#)
 - preserving configuration [7-6](#)
 - progression states [7-3](#)
 - reloading [7-5](#)

- reset [5-3](#)
- resetting [7-5](#)
- states [7-1](#)
- thresholds [8-9](#)
- switchover mechanism [5-6, 5-7](#)
 - HA [5-7, 7-3](#)
 - warm [7-3](#)
- switch priority
 - configuring [24-6](#)
 - range [2-25](#)
- switch redundancy states [5-6](#)
- switch states [12-6](#)
- syslogs
 - viewing [1-9](#)
- syslog server [26-2](#)
 - configuring [26-6](#)
- system assignment [22-53](#)
- system image [2-18](#)
 - reading configuration [6-26](#)
 - recovery interruption [6-27](#)
 - switching module [7-2](#)
- system images [6-1](#)
 - SYSTEM variable [6-2](#)
- system messages
 - configuring [26-4](#)
 - default settings [26-12](#)
 - displaying configuration [26-8](#)
 - logging [26-2](#)
- system processes
 - displaying [31-2](#)
 - status [31-5](#)
- system statistics
 - CPU and memory [31-6](#)
- system switchover
 - guidelines [5-3](#)
 - mechanisms [5-3](#)

T

- target disks [27-3](#)
- TCP connections
 - FCIP profiles [22-19](#)
- TCP ports
 - ACLs [20-7](#)
- TCP statistics
 - displaying [22-11](#)
- Telnet
 - default service [16-27](#)
 - session [6-17](#)
- Telnet session
 - message logging [26-4](#)
- temporary storage [2-17](#)
- TE port
 - trunking [1-6](#)
- TE ports
 - classes of service [10-4](#)
 - fctrace [29-4](#)
 - FSPF topology [19-2](#)
 - interface modes [10-2](#)
 - interoperability [29-22](#)
 - recovering from isolation [13-10](#)
 - SPAN [28-3](#)
 - trunking restrictions [11-1](#)
- TFTP
 - boot [6-29](#)
 - copying images [6-17](#)
 - server [6-29](#)
- TFTP server [31-6](#)
- threshold
 - major and minor [8-9](#)
- time interval
 - configuring [19-6](#)
- time out value
 - See TOV
- Timers
 - range [2-25](#)

TL Ports

- logging facility [26-3](#)

TL ports

- classes of service [10-4](#)
- configuring [10-10](#)
- displaying [10-27](#)
- FCS [30-2, 30-3](#)
- interface modes [10-2](#)
- SPAN [28-3](#)

TOV

- interoperability [29-21](#)
- ranges [29-2](#)

- transit VSAN [14-3](#)

- transit VSANs [14-7](#)

- guidelines [14-4](#)

- trivial authentication [16-2, 16-15](#)

troubleshooting

- error messages [26-2](#)

trunk-allowed list

- configuring [11-4](#)

Trunking

- PortChannels comparison [12-3](#)

trunking

- configuration guidelines [11-6](#)
- functionality [1-6](#)
- interoperability [29-21](#)
- link state [11-3](#)
- restrictions [11-1](#)

- trunking ports [9-5](#)

- trunking protocol [11-2, 11-6](#)

- default [11-2](#)
 - default settings [11-8](#)

trunk mode

- administrative default [10-14](#)
- configuring [11-3](#)
- default settings [11-8](#)
- status [11-3](#)

trunk ports

- displaying information [11-7](#)

U

UDP ports

- ACLs [20-7](#)

uninstalling

- permanent licenses [3-8](#)

updating

- licenses [3-9](#)

upgrades

- See disruptive upgrades
- See nondisruptive upgrades

upgrading

- software [6-17 to 6-23](#)

upgrading BIOS

- See BIOS upgrades

user ID

- authentication [16-3](#)

user IDs

- security management [1-10](#)

- user profiles [16-3](#)

users

- creating [16-32](#)

V

- VE ports [22-18](#)

- virtual devices [10-27](#)

virtual E ports

- See VE ports [22-18](#)

virtual Fibre Channel hosts

- mapping [22-52](#)

- virtual ISL [22-18](#)

Virtual LANs

- See VLANs

virtual N port

- dynamic mapping [22-52](#)

Virtual Router Redundancy Protocol

- See VRRP

virtual SANs

- See VSANs
- VLANs
 - configuring 22-5
- VLAN tags 22-6
- VR IDs
 - configuring 20-19
 - mapping 20-18
- VRRP
 - characteristics 20-18
 - clearing statistics 20-24
 - configuring 22-14
 - configuring Gigabit Ethernet 22-12
 - group members 22-13
 - logging facility 26-3
 - master and backup 20-18
 - primary IP 20-20
 - priority tracking 20-22
 - security authentication 20-22
 - setting priority 20-20
 - tracking priority 20-22
- VSA
 - communicating attributes 16-8
 - protocol options 16-8
- VSAN
 - address format 2-25
 - configuring 1-5
 - domain IDs 24-7
 - functionality 1-5
 - gateway switch 20-3
 - overlaid routes 20-4
 - reason codes 10-7
 - redundancy 1-5
 - scalability 1-5
 - traffic isolation 1-5
- VSAN IDs
 - allowed list 11-8
 - attributes 9-6
 - configuring FICON 21-3
 - FCS registration 2-8
 - membership 9-4
 - multiplexing traffic 10-4
 - name 9-10
 - range 9-5
 - trunking 12-3
- VSANs
 - allowed-active 11-1, 11-4
 - allowed list 28-4
 - allowed-list 11-8
 - attributes 9-6, 9-8
 - availability 9-1
 - broadcast address 19-10
 - cache contents 24-15
 - configuring 9-1, 9-6
 - configuring domains 24-1
 - configuring FSPF 19-4
 - configuring overlay 20-16
 - database submode 2-8
 - default setting 9-10
 - default VSAN 9-5
 - deleting 9-8
 - FCIDs 9-2
 - FCS 30-2
 - features 9-2
 - flow statistics 19-14
 - FSPF connectivity 19-2
 - functionality 1-5
 - interface 10-15, 10-17
 - interop mode 29-21
 - IP addresses 20-11
 - IPFC interface 29-4
 - isolated VSAN 9-5
 - logical interface 4-9
 - loop devices 10-27
 - management interfaces 20-2
 - membership 9-5, 9-9
 - merging traffic 11-6
 - mismatch 10-7, 11-2
 - multiple zones 9-4, 13-7

name [9-6](#)
 name server [15-3](#)
 overlaid routes [20-14](#)
 port granularity [9-3](#)
 port isolation [11-6](#)
 Rules and features [16-18](#)
 sate [9-6](#)
 scalability [9-1](#)
 SPAN source [28-2, 28-4](#)
 static routing [20-12](#)
 TOVs [29-2](#)
 traffic isolation [9-1, 9-3](#)
 traffic routing [20-1](#)
 trunk allowed [10-15](#)
 trunk-allowed [11-1, 11-2](#)
 trunk-allowed list [11-4](#)
 trunking port [10-4](#)
 trunking ports [9-5](#)
 usage [9-9](#)
 VRRP [20-18](#)
 VRRP submode [2-8](#)

VSAN trunking

See trunking

Z

zone database [13-11](#)
 zones
 access control [13-6](#)
 accesses between devices [1-5](#)
 configuring [13-4](#)
 configuring guidelines [13-7](#)
 default policy [13-2, 13-9](#)
 default settings [13-20](#)
 enforcing [13-9](#)
 examples [13-3](#)
 functionality [1-5](#)
 logging facility [26-3](#)
 See also default zones
 See also hard zoning
 See also soft zoning

W

world wide names

See WWNs

WWN

address pool [22-52](#)

WWNs

configuring [29-18](#)

displaying configurations [29-19](#)

suspended connection [10-8](#)

See also nWWNs

See also pWWNs